

NetBackup™ Administrator's Guide, Volume I

UNIX, Windows, and Linux

Release 11.1

NetBackup Administrator's Guide, Volume I

Last updated: 2025-11-24

Legal Notice

Copyright © 2025 Cohesity, Inc. All rights reserved.

Cohesity, Veritas, the Cohesity Logo, Veritas Logo, Veritas Alta, Cohesity Alta, and NetBackup are trademarks or registered trademarks of Cohesity, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Cohesity is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Cohesity product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Cohesity, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Cohesity, Inc. SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Cohesity as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Cohesity, Inc.
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Cohesity account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Cohesity website.

Cohesity Services and Operations Readiness Tools (SORT)

Cohesity Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	About NetBackup	32
Chapter 1	Introducing the NetBackup interfaces	33
	About NetBackup	33
	NetBackup documentation	35
	NetBackup administration interfaces	35
	About security certificates for NetBackup hosts	36
	About setting up the NetBackup Administration Console on UNIX	36
	Administering remote servers of different versions	37
	Logging in to the NetBackup Administration Console	38
	Using the NetBackup web UI	41
Section 2	Configuring hosts	45
Chapter 2	Configuring Host Properties	46
	About the NetBackup Host properties	48
	Methods to set the NetBackup configuration options	48
	Connecting to a host to view the host properties	49
	Changing the host properties on multiple hosts at the same time	50
	Exporting host properties	51
	Access Control properties	51
	Authentication Domain tab of the Access Control properties	52
	Authorization Service tab of the Access Control properties	53
	Network Attributes tab of the Access Control properties	53
	Active Directory properties	55
	Bandwidth properties	55
	Bandwidth limit usage considerations and restrictions	56
	Add Bandwidth Settings dialog box for Bandwidth properties	57
	Backup pool host properties	57
	Busy file settings properties	59
	Activating the Busy file settings in host properties	60
	Clean up properties	61

Client name properties	63
Client attributes properties	64
General tab of the Client attributes properties	66
Connect options tab of the Client attributes properties	70
Windows open file backup tab of the Client attributes properties	71
Client settings properties for UNIX clients	73
VxFS file change log (FCL) for incremental backups property	75
Client settings properties for Windows clients	77
How to determine if change journal support is useful in your NetBackup environment	80
Guidelines for enabling NetBackup change journal support	80
Cloud Storage properties	81
Credential access properties	82
Data Classification properties	83
Adding a Data Classification	84
Default job priorities properties	85
Understanding the job priority setting	86
Distributed application restore mapping properties	87
Encryption properties	88
Additional encryption methods for Windows clients	89
Enterprise Vault properties	90
Enterprise Vault hosts properties	91
Exchange properties	92
About the Exchange credentials in the client host properties	93
Exclude list properties	94
About the Add to Exclude List and Add Exceptions to Exclude List dialog boxes	95
Add an entry to an exclude list	96
Add an exception to the exclude list	96
Syntax rules for exclude lists	97
About creating an include list on a UNIX client	99
Traversing excluded directories	100
Fibre Transport properties	101
About Linux concurrent FT connections	103
Firewall properties	104
General server properties	106
Forcing restores to use a specific server	108
Global attributes properties	109
About constraints on the number of concurrent jobs	112
Setting up mailx email client	113
Logging properties	113

Logging levels	115
Lotus Notes properties	117
Media properties	119
Results when media overwrites are not permitted	122
Recommended use for Enable SCSI reserve property	123
Network properties	124
Network settings properties	124
Reverse host name lookup property	125
Use the IP address family property	126
Port ranges properties	127
Registered ports and dynamically-allocated ports	128
Preferred network properties	129
Add or edit a Preferred network setting	131
How NetBackup uses the directives to determine which network to use	133
Configurations to use IPv6 networks	136
Configurations to use IPv4 networks	138
Order of directive processing in the Preferred network properties	139
bptestnetconn utility to display Preferred network information	140
Configuration to prohibit using a specified address	141
Configuration to prefer a specified address	142
Configuration that restricts NetBackup to one set of addresses	143
Configuration that limits the addresses, but allows any interfaces	144
Properties setting in host properties	144
RHV access hosts properties	145
Resilient network properties	145
View the resiliency status of a client	147
About Resilient jobs	148
Resilient connection resource usage	148
Specify resilient connections for clients	149
Resource limit properties	150
Restore failover properties	151
Assigning an alternate media server as a failover restore server	152
Retention periods properties	152
Changing a retention period	154
Determining retention periods for volumes	155
Scalable Storage properties	155
Configuring advanced bandwidth throttling settings	157

Advanced bandwidth throttling settings	157
Servers properties	159
Adding a server to a servers list	160
Removing a server from a servers list	161
Enabling NetBackup clustered primary server inter-node authentication	161
About trusted primary servers	162
Changing the primary server that performs backups and restores for a client	166
SharePoint properties	167
Consistency check options for SharePoint Server	168
SLP settings properties	168
About batch creation logic in Storage Lifecycle Manager	173
Throttle bandwidth properties	174
Timeouts properties	175
Universal settings properties	178
User account settings properties	180
Terminate a NetBackup user session	180
Unlock a NetBackup user	181
Configure when idle sessions should time out	182
Configure the maximum of concurrent user sessions	182
Configure the maximum of failed sign-in attempts	183
Display a banner to users when they sign in	183
UNIX client properties	184
UNIX Server properties	184
VMware access hosts properties	185
Windows client properties	185
Configuration options not found in the host properties	186
About using commands to change the configuration options on UNIX or Linux clients and servers	186
Configuration options for NetBackup servers	187
ALLOW_MEDIA_OVERWRITE option for NetBackup servers	188
AUTO_ADD_ALL_ALIASES_FOR_CLIENT option for NetBackup servers	188
BPBRM_VERBOSE option for NetBackup servers	189
BPCD_ALLOWED_PATH option for NetBackup servers and clients	190
BPDBJOBS_COLDEFS options for Linux primary servers	191
BPDBM_VERBOSE option for NetBackup servers	195
BPRD_VERBOSE option for NetBackup servers	196
BPTM_VERBOSE option for NetBackup servers	197
BPEND_TIMEOUT option for NetBackup servers	199

BPSTART_TIMEOUT option for NetBackup servers	199
CHECK_RESTORE_CLIENT option for NetBackup servers	200
CLIENT_CONNECT_TIMEOUT option for NetBackup servers	201
CLIENT_PORT_WINDOW option for NetBackup servers and clients	201
CLIENT_READ_TIMEOUT option for NetBackup servers	202
CLOUD_AUTODISCOVERY_INTERVAL for NetBackup servers	204
CLUSTER_ECA_CERT_PATH for clustered primary server	204
CLUSTER_ECA_KEY_PASSPHRASEFILE for clustered primary server	205
CLUSTER_ECA_PRIVATE_KEY_PATH for clustered primary server	206
CLUSTER_ECA_TRUST_STORE_PATH for clustered primary server	206
COMPUTE_IMAGE_ENTROPY for NetBackup primary servers	207
CONNECT_OPTIONS option for NetBackup servers	208
DATAACCESS_AUDIT_INTERVAL_HOURS for NetBackup primary servers	210
DEFAULT_CONNECT_OPTIONS option for NetBackup servers	211
DISABLE_CERT_AUTO_RENEW option for NetBackup servers and clients	212
DISABLE_JOB_LOGGING option for NetBackup servers	213
DISABLE_STANDALONE_DRIVE_EXTENSIONS option for NetBackup servers	214
DISALLOW_BACKUPS_SPANNING_MEDIA option for NetBackup servers	214
DISALLOW_CLIENT_LIST_RESTORE option for NetBackup servers	215
DISALLOW_CLIENT_RESTORE option for NetBackup servers	216
DISALLOW_SERVER_FILE_WRITES option for NetBackup servers and clients	216
DTE_IGNORE_IMAGE_MODE for NetBackup servers	219
ECA_CERT_PATH for NetBackup servers and clients	220
ECA_CRL_CHECK for NetBackup servers and clients	223
ECA_CRL_PATH for NetBackup servers and clients	224
ECA_CRL_PATH_SYNC_HOURS for NetBackup servers and clients	226

ECA_CRL_REFRESH_HOURS for NetBackup servers and clients	226
ECA_DISABLE_AUTO_ENROLLMENT for NetBackup servers and clients	227
ECA_DR_BKUP_WIN_CERT_STORE for NetBackup servers and clients	228
ECA_KEY_PASSPHRASEFILE for NetBackup servers and clients	229
ECA_PRIVATE_KEY_PATH for NetBackup servers and clients	229
ECA_TRUST_STORE_PATH for NetBackup servers and clients	230
EAT_VERBOSE option for NetBackup servers and clients	232
ECA_WIN_CERT_STORE_TIME_LAG_MINUTES for NetBackup servers and clients	233
ECMS_HOSTS_SECURE_CONNECT_ENABLED for servers	234
ENABLE_CRITICAL_PROCESS_LOGGING for NetBackup servers and clients	235
ENABLE_DIRECT_CONNECTION for servers	236
ENABLE_NBSQLADM option for NetBackup servers and clients	237
FAILOVER_RESTORE_MEDIA_SERVERS option for NetBackup servers	237
FETCH_MSDP_EKMS_KEY for NetBackup servers and clients	238
FORCE_RESTORE_MEDIA_SERVER option for NetBackup servers	239
GENERATE_ENGLISH_LOGS option for NetBackup servers and clients	240
GUI_ACCOUNT_LOCKOUT_DURATION option for NetBackup servers	241
GUI_IDLE_TIMEOUT option for NetBackup servers	241
GUI_MAX_CONCURRENT_SESSIONS option for NetBackup servers	242
GUI_MAX_LOGIN_ATTEMPTS option for NetBackup servers	243
HOSTDB_RESYNC_INTERVAL option for NetBackup servers and clients	243
HYPERV_WMI_CREATE_DISK_TIMEOUT option for NetBackup servers	244
INCOMPLETE_JOB_CLEAN_INTERVAL option for NetBackup servers and clients	245

INITIAL_BROWSE_SEARCH_LIMIT option for NetBackup servers and clients	246
INITIATE_REVERSE_CONNECTION for servers	247
IP_ADDRESS_FAMILY option for NetBackup servers	248
JOB_PRIORITY option for NetBackup servers and clients	249
KEEP_LOGS_SIZE_GB for NetBackup servers and clients	251
KMS_CONFIG_IN_CATALOG_BKUP for NetBackup primary server	252
LIMIT_BANDWIDTH option for NetBackup servers	253
MALWARE_DETECTION_JOBS_PER_SCAN_HOST option for NetBackup servers	253
MALWARE_SCAN_OPERATION_TIMEOUT	254
MANAGE_WIN_CERT_STORE_PRIVATE_KEY option for NetBackup primary servers	255
MAX_LOGFILE_SIZE option for NetBackup servers and clients for legacy logging	256
MAX_NUM_LOGFILES option for NetBackup servers and clients for legacy logging	257
MEDIA_UNMOUNT_DELAY option for NetBackup servers	258
MEDIA_REQUEST_DELAY option for NetBackup servers	259
MEDIA_SERVER option for NetBackup servers	260
MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB option for NetBackup servers	260
MPX_RESTORE_DELAY option for NetBackup servers	261
MUST_USE_LOCAL_DRIVE option for NetBackup servers	262
NAT_SERVER_LIST for servers	263
NB_FIPS_MODE option for NetBackup servers and clients	263
NBFP_CHECK option for NetBackup servers and clients	264
NBRNTD_IDLE_TIMEOUT option for NetBackup servers	264
NBSD_POLL_INTERVAL option for NetBackup servers and clients	265
NBSD_DUMP_COUNT option for NetBackup servers and clients	266
NBSD_MONITOR_CPU option for NetBackup servers and clients	267
NBSD_MONITOR_MEMORY option for NetBackup servers and clients	267
NBSD_MEMORY_UNIT option for NetBackup servers and clients	268
NBSD_MONITOR_DEADLOCK option for NetBackup servers and clients	269
NBSD_DEADLOCK_INTERVAL option for NetBackup servers and clients	269

NBSD_ALWAYS_DUMP option for NetBackup servers and clients	270
NBSD_CAPTURE_PROCESS_DUMP option for NetBackup servers and clients	271
NBSD_INCREASE_LOG_LEVEL option for NetBackup servers and clients	271
NBSD_CAPTURE_NETWORK_STAT option for NetBackup servers and clients	272
NBSD_CAPTURE_DISK_IO option for NetBackup servers and clients	272
NBSD_NUMBER_OF_READINGS option for NetBackup servers and clients	273
NBSD_READING_INTERVAL option for NetBackup servers and clients	274
NBSD_PURGE_OLD_EVIDENCE option for NetBackup servers and clients	274
NBSD_CAPTURE_WITHOUT_THRESHOLD option for NetBackup servers and clients	275
NBSD_JDK_HOME option for NetBackup servers and clients	276
NBSD_EVIDENCE_PATH option for NetBackup servers and clients	276
NBSD_VERBOSE option for NetBackup servers and clients	277
NBSD_AUTO_MONITOR option for NetBackup servers and clients	277
NBSD_AUTOMONITOR_CPU_THRESHOLD option for NetBackup servers and clients	278
NBSD_AUTOMONITOR_MEMORY_THRESHOLD option for NetBackup servers and clients	279
NBSD_MONITOR_POLICY_NAME option for primary server	280
NBSD_MONITOR_SYSTEM_FOR_HOURS option for NetBackup servers and clients	281
NBSD_EVIDENCE_SIZE_LIMIT option for NetBackup servers and clients	281
NBSD_PUSH_MONITOR_DATA_TO_REMOTE option for NetBackup servers and clients	282
NETBACKUP_NATIVE_AUDITING option for NetBackup primary server	283
NOTIFY_SNOOZE_PERIOD_IN_DAYS option for NetBackup servers	283
PREFERRED_NETWORK option for NetBackup servers	285

RANDOM_PORTS option for NetBackup servers and clients	298
RE_READ_INTERVAL option for NetBackup servers	299
REQUIRED_NETWORK option for NetBackup servers	299
RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH option for NetBackup servers	300
RESILIENT_BACKUP_JOB_RESTART_TIMEOUT option for NetBackup servers	301
RESILIENT_NETWORK option for NetBackup primary servers and clients	302
RESILIENT_RECONNECT_TIMEOUT	303
RESILIENT_RETRY_INTERVAL	304
RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers	304
REVERSE_NAME_LOOKUP option for NetBackup servers and clients	305
SECURE_PROXY_CIPHER_LIST option for NetBackup servers and clients	306
SERVER option for NetBackup servers	307
SERVER_CONNECT_TIMEOUT option for NetBackup servers	309
SERVER_PORT_WINDOW option for NetBackup servers	309
SERVER_RESERVED_PORT_WINDOW option for NetBackup servers and clients	310
SKIP_RESTORE_TO_SYMLINK_DIR option for NetBackup servers	311
SYSLOG_AUDIT_CATEGORIES for NetBackup primary server	312
SYSLOG_AUDIT_USE_OCSF_FORMAT for NetBackup primary server	313
TELEMETRY_UPLOAD option for NetBackup servers	314
THROTTLE_BANDWIDTH option for NetBackup servers	315
TRUSTED_PRIMARY option for NetBackup servers	318
ULINK_ON_OVERWRITE option for NetBackup servers	318
USE_URANDOM for NetBackup servers and clients	320
VERBOSE option for NetBackup servers and clients	321
VIRTUALIZATION_CRL_CHECK for NetBackup servers and clients	322
VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED for servers and clients	323
VIRTUALIZATION_HOSTS_CONNECT_TIMEOUT for servers and clients	324

VMWARE_AUTODISCOVERY_INTERVAL option for NetBackup servers	325
VM_SNAPSHOT_QUIESCE_STATUS option for NetBackup servers and clients	326
VMWARE_TLS_MINIMUM_V1_2 for NetBackup servers and clients	327
WEB_SERVER_TUNNEL_ENABLE option for NetBackup servers	328
Configuration options for NetBackup clients	329
ACCEPT_REVERSE_CONNECTION for clients	329
APP_PROXY_SERVER option for NetBackup clients	330
BACKUP_BTRFS_SNAPSHOT option for NetBackup clients	330
BACKUP_FIFO_FILES option for NetBackup clients	331
BPARCHIVE_POLICY option for NetBackup clients	332
BPARCHIVE_SCHED option for NetBackup clients	333
BPBACKUP_POLICY option for NetBackup clients	334
BPBACKUP_SCHED option for NetBackup clients	335
BUSY_FILE_ACTION option for NetBackup clients	336
BUSY_FILE_DIRECTORY option for NetBackup clients	337
BUSY_FILE_NOTIFY_USER option for NetBackup clients	338
BUSY_FILE_PROCESSING option for NetBackup clients	339
CLIENT_NAME option for NetBackup clients	339
COMPRESS_SUFFIX option for NetBackup clients	340
CRYPT_CIPHER option for NetBackup clients	341
CRYPT_KIND option for NetBackup clients	342
CRYPT_OPTION option for NetBackup clients	343
CRYPT_STRENGTH option for NetBackup clients	344
CRYPT_LIBPATH option for NetBackup clients	345
CRYPT_KEYFILE option for NetBackup clients	346
DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients	347
DTE_CLIENT_MODE for clients	348
ENABLE_DATA_CHANNEL_ENCRYPTION for clients	349
IGNORE_XATTR option for NetBackup clients	350
INFORMIX_HOME option for NetBackup clients	353
KEEP_DATABASE_COMM_FILE option for NetBackup clients	353
KEEP_LOGS_DAYS option for NetBackup clients	354
LIST_FILES_TIMEOUT option for NetBackup clients	355
LOCKED_FILE_ACTION option for NetBackup clients	356
MEDIA_SERVER option for NetBackup clients	356
MEGABYTES_OF_MEMORY option for NetBackup clients	357
MSSQL_CONFIG_LIST for SQL Server clients	358

	MSSQL_ODBC_ENCRYPT_CONNECTION for SQL Server clients	359
	MSSQL_ODBC_PREFERRED_DRIVER for SQL Server clients	360
	MSSQL_ODBC_TRUST_SERVER_CERTIFICATE for SQL Server clients	361
	NFS_ACCESS_TIMEOUT option for NetBackup clients	362
	OLD_VNETD_CALLBACK option for NetBackup clients	362
	REPORT_CLIENT_DISCOVERIES option for NetBackup clients	363
	RESTORE_RETRIES option for NetBackup clients	364
	RMAN_OUTPUT_DIR for NetBackup clients	365
	SERVER option for NetBackup clients	366
	SUBSCRIBER_HEARTBEAT_TIMEOUT for clients	367
	SYBASE_HOME option for NetBackup clients	368
	USE_CTIME_FOR_INCREMENTALS option for NetBackup clients	369
	USE_FILE_CHG_LOG option for NetBackup clients	369
	USEMAIL option for NetBackup clients	370
	WEB_SERVER_TUNNEL option for NetBackup clients	371
	WEB_SERVER_TUNNEL_USE option for NetBackup clients	372
Chapter 3	Configuring server groups	374
	About NetBackup server groups	374
	Add a server group	374
	Delete a server group	375
Chapter 4	Enabling support for NAT clients and NAT servers in NetBackup	376
	About NAT support in NetBackup	376
	Important notes	378
	Workflow to enable NAT hosts in NetBackup domain	379
	Configuring the NetBackup Messaging Broker service	380
	Configuring NetBackup Messaging Broker service to use PBX port	381
	Removing NAT support from NetBackup	382
	Communication with clients other than NAT clients	382
	Performance characteristics of NAT support	383

Chapter 5	Configuring host credentials	384
	About configuring credentials	384
	About configuring Snapshot Management server in NetBackup	385
	Registering a Snapshot Manager server in NetBackup	385
	Configuring Snapshot Manager plug-ins in NetBackup	387
Chapter 6	Managing media servers	389
	Activating or deactivating a media server	389
	Adding a media server	390
	Registering a media server	391
	Deleting all devices from a media server	392
	Removing a device host from the NetBackup database	395
	About decommissioning a media server	395
	About decommissioning limitations	396
	Before you decommission a media server	397
	Post decommission recommendations	397
	Decommission actions	398
	Previewing references to a media server	401
	Decommissioning a media server	402
Section 3	Configuring storage	405
Chapter 7	Configuring disk storage	406
	About configuring BasicDisk storage	406
	About configuring disk pool storage	406
	Configuring NetBackup MSDP disk pools	407
	About disk pools for NetBackup deduplication	407
	Configuring a disk pool for deduplication	408
	Managing Media Server Deduplication Pools	410
Chapter 8	Configuring robots and tape drives	417
	NetBackup robot types	418
	About the device mapping files	419
	Downloading the device mapping files	419
	About configuring robots and tape drives in NetBackup	420
	About device discovery	421
	About device serialization	422
	About adding devices without discovery	423
	About robot control	423
	About drive name rules	424

Configuring robots and tape drives by using the wizard	425
Updating the device configuration by using the wizard	425
Adding a robot to NetBackup manually	426
Robot properties and configuration options	428
Managing robots	430
Changing robot properties	430
Delete a robot	431
Moving a robot and its media to a new media server	431
Adding a tape drive to NetBackup manually	432
Tape drive configuration options	434
Configure drive name rules	437
Adding a tape drive path	438
About SCSI reserve on drive paths	439
Drive path options	439
About no rewind device files on UNIX	440
Adding a shared tape drive to a NetBackup environment	441
Correlating tape drives and SCSI addresses on Windows hosts	441
Correlating tape drives and device files on UNIX hosts	443
UNIX device correlation example	444
Managing tape drives	444
Changing a drive comment	445
About downed drives	445
Changing a drive operating mode	446
Changing a tape drive path	446
Drive path options	447
Changing the operating mode for a drive path	448
Changing tape drive properties	449
Changing a tape drive to a shared drive	449
Cleaning a tape drive from the Device monitor	450
Deleting a drive	450
Resetting a drive	451
Resetting the mount time of a drive	451
Setting the drive cleaning frequency	452
Viewing drive details	452
Performing device diagnostics	453
Running a robot diagnostic test	453
Running a tape drive diagnostic test	454
Managing a diagnostic test step that requires operator intervention	455
Obtaining detailed information for a diagnostic test step	455
Verifying the device configuration	456
About automatic path correction	456
Enabling automatic path correction	457

Replacing a device	457
Updating device firmware	459
About the NetBackup Device Manager	460
About external access to NetBackup controlled devices on UNIX	460
Stopping and restarting the device manager	461

Chapter 9 Configuring tape media 462

About NetBackup tape volumes	463
About NetBackup volume pools	463
About reserved volume pool name prefixes	464
About scratch volume pools	465
About NetBackup volume groups	466
NetBackup media types	466
About WORM media	468
About using volume pools to manage WORM media	469
About using unique drive and media types to manage WORM media	470
Disabling WORM volume pool name verification	471
About WORM media and the Quantum drive	471
Supported WORM drives	472
About adding volumes	472
About adding robotic volumes	472
About adding standalone volumes	473
About configuring media name and attribute rules	473
Adding volumes	474
Volume properties	475
About media settings rules	478
Configuring media settings	478
Media settings options	479
About barcodes	486
About barcode rules	488
Configuring barcode rules	490
Barcode rules settings	491
About media ID generation rules	492
Configuring media ID generation rules	493
Media ID generation options	495
About media type mapping rules	496
Configuring media type mappings	496
About adding media type mapping entries	499
Default and allowable media types	499
Managing volumes	503
Changing the group of a volume	503

About rules for moving volumes between groups	504
Changing the owner of a volume	504
Edit a volume	505
About assigning and deassigning volumes	505
Delete a volume	506
Erasing a volume	507
About exchanging a volume	508
About frozen media	510
About injecting and ejecting volumes	511
About rescanning and updating barcodes	513
About labeling NetBackup volumes	515
About moving volumes	517
About recycling a volume	519
Suspending or unsuspending volumes	520
Managing volume pools	521
Add a volume pool	521
Edit or delete a volume pool	521
Volume pool properties	522
Managing volume groups	524
Moving a volume group	524
Deleting a volume group	524
About media sharing	525
Configure unrestricted media sharing	526
Configure media sharing with a server group	526

Chapter 10	Inventorying robots	528
	About robot inventory	528
	When to inventory a robot	529
	About showing a robot's contents	531
	About inventory results for API robots	532
	Showing the media in a robot	533
	About comparing a robot's contents with the volume configuration	534
	Comparing media in a robot with the volume configuration	535
	About previewing volume configuration changes	535
	Previewing volume configuration changes for a robot	536
	About updating the NetBackup volume configuration	537
	Volume update prerequisites	538
	Updating the NetBackup volume configuration with a robot's contents	538
	Robot inventory options	539
	About the vmphyinv physical inventory utility	540

	How vmphynv performs a physical inventory	542
Chapter 11	Configuring storage units	547
	About storage	547
	Creating a storage unit	548
	Creating a storage unit by copying a storage unit	550
	Editing storage unit settings	551
	Deleting storage units	552
	Tape storage unit considerations	553
	Disk storage unit considerations	554
	NDMP storage unit considerations	558
	About storage unit settings	559
	Absolute pathname to directory or absolute pathname to volume setting for storage units	559
	Density storage unit setting	560
	Disk type storage unit setting	560
	Enable block sharing storage unit setting	561
	Enable multiplexing storage unit setting	561
	High water mark storage unit setting	561
	Low water mark storage unit setting	562
	Maximum concurrent write drives storage unit setting	563
	Maximum concurrent jobs storage unit setting	563
	Maximum streams per drive storage unit setting	565
	Media server storage unit setting	565
	NDMP host storage unit setting	568
	On demand only storage unit setting	568
	Only use the following media servers storage unit setting	568
	Properties option in the Change Storage Units dialog box	569
	Reduce fragment size storage unit setting	571
	Robot number storage unit setting	572
	Robot type storage unit setting	572
	Select disk pool storage unit setting	572
	Staging schedule option in Change Storage Units dialog	572
	Storage device setting for storage units	573
	Storage unit name setting	573
	Storage unit type setting	573
	Enable temporary staging area storage unit setting	573
	Use any available media server storage unit setting	574
	Use WORM setting	574
	Overview of universal shares	575

Chapter 12	Staging backups	576
	About staging backups	576
	About basic disk staging	577
	Creating a basic disk staging storage unit	578
	Creating a schedule for a BasicDisk staging storage unit	580
	Configuring multiple copies in a relocation schedule	581
	Disk staging storage unit size and capacity	582
	Finding potential free space on a BasicDisk disk staging storage unit	584
	Disk Staging Schedule dialog box	585
	Initiating a relocation schedule manually	588
Chapter 13	Configuring storage unit groups	590
	About storage unit groups	590
	Creating storage unit groups for backups	591
	Creating storage unit groups for snapshots	593
	Deleting a storage unit group	595
	Storage unit selection criteria within a group	595
	Media server load balancing	596
	Exception to the storage unit selection criteria	598
	About disk spanning within storage unit groups	599
Section 4	Configuring storage lifecycle policies (SLPs)	600
Chapter 14	Configuring storage lifecycle policies	601
	About storage lifecycle policies	601
	Creating a storage lifecycle policy	602
	Modifying the hierarchy of operations in a storage lifecycle policy	603
	Deleting a storage lifecycle policy	603
	Lifecycle operation administration using the nbstlutil command	605
Chapter 15	Storage operations	607
	Operation types in a storage lifecycle policy	607
	Backup operation in an SLP	608
	Backup From Snapshot operation in an SLP	609
	Duplication operation in an SLP	610
	Import operation in an SLP	611

Index From Snapshot operation in an SLP	612
Determining where and when the Index From Snapshot operation occurs	614
Replication operation in an SLP	615
Snapshot operation in an SLP	616
Primary snapshot storage unit	618
Primary + Replication source snapshot storage unit	618
Replication source + Replication target snapshot storage unit	619
Replication target snapshot storage unit	619
Replication source + Replication target + Mirror snapshot storage unit	620
Replication target + Mirror snapshot storage unit	620
Creating a hierarchy of storage operations in a storage lifecycle policy	620

Chapter 16 Retention types for SLP operations 622

Retention types for storage lifecycle policy operations	622
Capacity managed retention type for SLP operations	624
Rules and recommendations for using the Capacity Managed retention type	625
Capacity managed retention type and disk types that support SIS	625
Expire after copy retention type for SLP operations	626
Fixed retention type for SLP operations	626
Maximum snapshot limit retention type for SLP operations	627
Mirror retention type for SLP operations	627
Target retention type for SLP operations	628

Chapter 17 Storage lifecycle policy options 629

Storage lifecycle policy settings	629
SLP settings	631
Properties tab for an SLP operation	632
Window tab	634
Creating a new window for a storage lifecycle policy operation	635
Excluding days from a window for a storage lifecycle policy operation	636
Storage lifecycle policy Validation report tab	637

Chapter 18	Using a storage lifecycle policy to create multiple copies	639
	About writing multiple copies using a storage lifecycle policy	639
	How the order of the operations determines the copy order	640
	About ensuring successful copies using lifecycles	640
Chapter 19	Storage lifecycle policy versions	642
	About storage lifecycle policy versions	642
	Storage lifecycle changes and versioning	643
	When changes to storage lifecycle policies become effective	644
	Deleting old storage lifecycle policy versions	645
Section 5	Configuring backups	647
Chapter 20	Creating backup policies	648
	About the Policies utility	649
	Planning for policies	650
	Windows example of one client in multiple policies	652
	Policy attributes that affect how clients are grouped in policies	653
	About Microsoft DFSR backups and restores	654
	Creating a backup policy	657
	Adding or changing schedules in a policy	658
	Changing multiple policies at one time	659
	About automanaged policies or storage lifecycle policies	660
	Copying or moving policy items to another policy	660
	Copying a policy to create a new policy	661
	Copying a schedule into the same policy or different policy	662
	Deleting schedules, backup selections, or clients from a policy	662
	Policy Attributes tab	663
	Policy type (policy attribute)	664
	Data classifications (policy attribute)	668
	Policy storage (policy attribute)	669
	Policy volume pool (policy attribute)	672
	Take checkpoints every __ minutes (policy attribute)	674
	Limit jobs per policy (policy attribute)	678
	Job priority (policy attribute)	680
	Media Owner (policy attribute)	681
	Go into effect at (policy attribute)	681
	Follow NFS (policy attribute)	682

Backup Network Drives (policy attribute)	683
Cross mount points (policy attribute)	685
Compression (policy attribute)	689
Encryption (policy attribute)	691
Collect disaster recovery information for Bare Metal Restore (policy attribute)	693
Collect true image restore information (policy attribute) with and without move detection	693
Allow multiple data streams (policy attribute)	697
Use multiple MSDP nodes (policy attribute)	701
Client-side deduplication (policy attribute)	701
Enable granular recovery (policy attribute)	702
Use Accelerator (policy attribute)	703
Enable optimized backup of Windows deduplicated volumes	723
Keyword phrase (policy attribute)	727
Snapshot Client and Replication Director (policy attributes)	728
Perform block level incremental backups (policy attributes)	728
Use Replication Director (policy attributes)	728
Perform snapshot backups (policy attributes)	730
Microsoft Exchange Attributes (policy attributes)	730
Schedules tab	730
Schedule Attributes tab	731
Name (schedule attribute)	732
Type of backup (schedule attribute)	732
Synthetic backup (schedule attribute)	741
Accelerator forced rescan option (schedule attribute)	742
Calendar (schedule attribute)	744
Frequency (schedule attribute)	744
Instant Recovery (schedule attribute)	746
Multiple copies (schedule attribute)	747
Override policy storage (schedule attribute)	751
Override policy volume pool (schedule attribute)	752
Override media owner (schedule attribute)	752
Retention (schedule attribute)	753
Media multiplexing (schedule attribute)	756
Start Window tab	763
Adding, changing, or deleting a time window in a policy schedule	763
Example of schedule duration	765
Excluding days from a schedule	766
Include Dates tab	767
Calendar scheduling with the Include Dates tab	768
How NetBackup determines which schedule to run next	769

About schedule windows that span midnight	772
How open schedules affect calendar-based and frequency-based schedules	773
Creating an open schedule in the NetBackup Administration Console	777
Runtime considerations that affect backup frequency	778
About the Clients tab	779
Adding, changing, or deleting clients in a policy	779
Browse for Hyper-V virtual machines	782
Backup Selections tab	782
Adding backup selections to a policy	783
Verifying the Backup Selections list	791
How to reduce backup time	794
Pathname rules for Windows client backups	795
Pathname rules for Windows disk image (raw) backups	797
Pathname rules for Windows registry backups	799
About hard links to files and directories	800
Pathname rules for UNIX client backups	802
Pathname rules for the clients that run extension products	810
About the directives on the Backup Selections list	810
Files that are excluded from backups by default	823
About excluding files from automatic backups	825
Files that are excluded by Microsoft Windows Backup	826
Disaster Recovery tab	827
Adding policies to the Critical Policies list of a catalog backup policy	830
Creating a Vault policy	830
Creating a BigData policy	832
Performing manual backups	834
Active Directory granular backups and recovery	835
System requirements for Active Directory granular NetBackup backups and recovery	836
Creating a policy that allows Active Directory granular restores	836
Restoring Active Directory objects	838
Troubleshooting granular restore issues	839

Chapter 21	Synthetic backups	842
	About synthetic backups	842
	Recommendations for synthetic backups and restores	843
	Synthetic full backups	846
	Synthetic cumulative incremental backups	848

Schedules that must appear in a policy for synthetic backups	850
Adding clients to a policy for synthetic backups	850
Change journal and synthesized backups	851
True image restore and synthesized backups	851
Displaying synthetic backups in the Activity Monitor	851
Logs produced during synthetic backups	852
Synthetic backups and directory and file attributes	852
Using the multiple copy synthetic backups method	853
Configuring multiple copy synthetic backups	854
Configuration variables for multiple copy synthetic backups	855
Multiple copy synthetic backups configuration examples	856
Optimized synthetic backups	857
Optimized synthetic backups for deduplication	857

Chapter 22	Protecting the NetBackup catalog	858
	About the NetBackup catalog	858
	Parts of the NetBackup catalog	859
	NetBackup databases and configuration files	860
	About the NetBackup image database	862
	About the catalog backup of cloud configuration files	864
	Catalog backups	865
	The catalog backup process	865
	Strategies that ensure successful NetBackup catalog backups	866
	Prerequisites for backing up the NetBackup catalog	867
	Configuring catalog backups	867
	Backing up NetBackup catalogs manually	869
	Concurrently running catalog backups with other backups	870
	Catalog policy schedule considerations	870
	How catalog incrementals and standard backups interact on UNIX	871
	Determining whether or not a catalog backup succeeded	871
	Recovering the catalog	872
	Disaster recovery emails and the disaster recovery files	872
	Disaster recovery packages	873
	About disaster recovery settings	874
	Setting a passphrase to encrypt disaster recovery packages	875
	Archiving the catalog and restoring from the catalog archive	877
	Enabling intelligent catalog archiving (ICA) to reduce the number	
	of .f files	880
	Creating a catalog archiving policy	884
	Catalog archiving commands	885

	Catalog archiving considerations	887
	Extracting images from the catalog archives	888
	Estimating catalog space requirements	888
	NetBackup file size considerations on UNIX systems	890
	Moving the image catalog	890
	About image catalog compression	892
Chapter 23	About the NetBackup database	896
	About the NetBackup database installation	896
	About NetBackup primary server installed directories and files	896
		896
	NetBackup configuration entry	899
	NetBackup database server management	900
	The NetBackup database and clustered environments	901
	Post-installation tasks	901
	Changing the NetBackup database password	902
	Moving a database after installation	903
	Copying the NetBackup databases	905
	Creating the NBDB database manually	905
	Using the NetBackup Database Administration utility on Windows	907
	General tab of the NetBackup Database Administration utility	909
		909
	Tools tab of the NetBackup Database Administration utility	911
	Using the NetBackup Database Administration utility on UNIX	915
	Select/Restart Database and Change Password menu options	916
		916
	Database Space Management menu options	917
	Database Validation Check and Rebuild menu options	918
	Move Database menu options	919
	Unload Database menu options	920
	Backup and Restore Database menu options	920
Chapter 24	Managing backup images	922
	About the Catalog utility	922
	Catalog utility search criteria and backup image details	923
	Verifying backup images	926
	Promoting a copy to a primary copy	927
	Duplicating backup images	928
	Multiplexed duplication considerations	932
	Jobs that appear while making multiple copies	932
	Expiring backup images	933

	About Image Dependency Expiration Cleanup	933
	About importing backup images	935
	About importing expired images	935
	Importing backup images, Phase I	936
	Importing backup images, Phase II	937
Chapter 25	Configuring immutability and indelibility of data in NetBackup	939
	About immutable and indelible data	939
	Workflow to configure immutable and indelible data	941
	Deleting an immutable image from storage using the <code>bpexpdate</code> command	942
	Removing an immutable image from the catalog using the <code>bpexpdate</code> command	944
Section 6	Deployment Management	945
Chapter 26	Deployment Management	946
	About the deployment policies utility	946
	Deployment policy management	947
	Copying a deployment policy to create a new deployment policy	951
	Attributes tab in Deployment management	952
	Schedules tab in Deployment management	953
	Security options tab in Deployment management	954
	Manually deploy a deployment policy	956
	Perform client initiated upgrade with VxUpdate	956
	Deployment job status	957
Section 7	Configuring replication	959
Chapter 27	About NetBackup replication	960
	About NetBackup replication	960
	About NetBackup Auto Image Replication	961
	One-to-many Auto Image Replication model	963
	Cascading Auto Image Replication model	963
	About the domain relationship for replication	966
	About the replication topology for Auto Image Replication	967
	Viewing the replication topology for Auto Image Replication	968
	About trusted primary servers for Auto Image Replication	974

About the storage lifecycle policies required for Auto Image Replication	977
About Auto Image Replication import confirmation	981
Auto Image Replication setup overview	982
How to resolve volume changes for Auto Image Replication	982
Removing or replacing replication relationships in an Auto Image Replication configuration	986
About restoring from a backup at a target primary domain	1000
Reporting on Auto Image Replication jobs	1001
About NetBackup Replication Director	1002

Section 8 Monitoring and reporting 1004

Chapter 28 Monitoring NetBackup activity 1005

About the Activity Monitor	1005
Setting Activity Monitor options	1007
About the Jobs tab	1009
Viewing job details in the Activity Monitor	1010
Deleting completed jobs in the Activity Monitor	1010
Canceling a job that has not completed in the Activity Monitor	1011
Restarting a failed (completed) job in the Activity Monitor	1011
Suspending and resuming jobs in the Activity Monitor	1011
Changing the Job Priority dynamically from the Activity Monitor	1012
About the Daemons tab	1013
Using the nbrbutil utility to configure the NetBackup Resource Broker	1017
Types of NetBackup daemons	1022
Monitoring NetBackup daemons	1022
Starting or stopping a daemon	1023
Displaying all media servers in the Activity Monitor	1023
About the Processes tab	1024
Monitoring NetBackup processes in the Process Details dialog box	1029
About the Drives tab	1030
Monitoring tape drives	1030
Cleaning tape drives from the Activity Monitor	1031
About the Error Logs tab	1031
About the jobs database	1032
Changing the default bpdjobs_options values	1033
About the BPDBJOBS_OPTIONS environment variable	1034

	bpdjobs command line options	1036
	Enabling the bpdjobs debug log	1036
	About the Device Monitor	1037
	About media mount errors	1037
	About pending requests and actions	1038
	About pending requests for storage units	1039
	Resolving a pending request	1039
	Resolving a pending action	1040
	Resubmitting a pending request	1041
	Denying a pending request	1041
Chapter 29	Reporting in NetBackup	1042
	About the Reports utility	1042
	Running a report	1044
	Copying report text to another document	1045
	Saving or exporting a report	1046
	Printing a report	1046
Chapter 30	Email notifications	1047
	Send notifications to the backup administrator about failed backups	1047
	Send notifications to a host administrator about backups	1048
	Configure the nbmail.cmd script on the Windows hosts	1048
	Send notifications about KMS certificate expiration	1050
Section 9	Administering NetBackup	1051
Chapter 31	Management topics	1052
	Configuring the NetBackup Client Service	1052
	Units of measure used with NetBackup	1053
	NetBackup naming conventions	1054
	Wildcard use in NetBackup	1055
Chapter 32	Accessing a remote server	1058
	Prerequisites for accessing a remote server	1058
	Allow access to another server	1058
	Authorize users of one server to access another server	1059
	Accessing remote servers	1060
	Troubleshooting remote server administration	1061

Chapter 33	Using the NetBackup Remote Administration Console	1063
	About the NetBackup Remote Administration Console	1063
	Run-time configuration options for the NetBackup Administration Console	1065
	BROWSER_BINARY_PATH	1065
	DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT	1066
	FIREWALL_IN	1066
	FORCE_IPADDR_LOOKUP	1068
	INITIAL_MEMORY, MAX_MEMORY	1069
	MEM_USE_WARNING	1070
	NB_FIPS_MODE	1070
	NBJAVA_CLIENT_PORT_WINDOW	1070
	NBJAVA_CORBA_DEFAULT_TIMEOUT	1071
	NBJAVA_CORBA_LONG_TIMEOUT	1071
	NETBACKUP_API_CLIENT_CONNECTION_TIMEOUT	1072
	NETBACKUP_API_CLIENT_READ_TIMEOUT	1072
	PBX_PORT	1072
	USE_URANDOM	1072
	VNETD_PORT	1073
	About improving NetBackup performance	1073
	About running the NetBackup Administration Console locally	1074
	About running a console locally and administering a remote server	1074
	Enhancing console performance	1075
	Determining better performance when the console is run locally or uses remote display back	1076
	About adjusting time zones in the NetBackup Administration console	1077
	Adjusting the time zone in the NetBackup Administration Console or the Backup, Archive, and Restore console	1078
	Configuring a custom time zone in the NetBackup Administration Console or the Backup, Archive, and Restore console	1078
	Time zone table	1079
Chapter 34	Alternate server restores	1103
	About alternate server restores	1103
	About supported configurations for alternate server restores	1104
	About performing alternate server restores	1105
	About modifying the NetBackup catalogs	1106
	Overriding the original server for restores	1107

	About enabling automatic failover to an alternate server	1109
	Expiring and importing media for alternate server restores	1110
Chapter 35	Managing client backups and restores	1112
	About server-directed restores	1112
	About client-redirected restores	1114
	About restore restrictions	1115
	Allowing all clients to perform redirected restores	1115
	Allowing a single client to perform redirected restores	1116
	Allowing redirected restores of a specific client's files	1117
	Examples of redirected restores	1117
	About restoring the files that have Access Control Lists (ACLs)	1123
	About setting the original atime for files during restores on UNIX	1125
	Restoring the System State	1125
	About the backup and restore of compressed files on VxFS file systems	1127
	About backups and restores on ReFS	1128
Chapter 36	Powering down and rebooting NetBackup servers	1129
	Shut down and restart NetBackup servers	1129
	Shut down and start up all NetBackup services and daemons	1130
	Restart a NetBackup server	1131
	Restart a NetBackup media server	1131
Chapter 37	About Granular Recovery Technology	1133
	About installing and configuring Network File System (NFS) for Active Directory Granular Recovery	1133
	About configuring Services for Network File System (NFS)	1134
	Enabling Services for Network File System (NFS) on a media server	1135
	Enabling Services for Network File System (NFS) on a client	1138
	Disabling the Client for NFS on the media server	1140
	Disabling the Server for NFS	1142
	Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology (GRT)	1144
	Configuring a different network port for NBFSD	1145

About NetBackup

- [Chapter 1. Introducing the NetBackup interfaces](#)

Introducing the NetBackup interfaces

This chapter includes the following topics:

- [About NetBackup](#)
- [NetBackup documentation](#)
- [NetBackup administration interfaces](#)
- [Using the NetBackup web UI](#)

About NetBackup

NetBackup provides a complete, flexible data protection solution for a variety of platforms. The platforms include Windows, UNIX, and Linux systems.

NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. An administrator can carefully schedule backups to achieve systematic and complete backups over a period of time, and optimize network traffic during off-peak hours. The backups can be full or incremental: Full backups back up all indicated client files, while incremental backups back up only the files that have changed since the last backup.

The NetBackup administrator can allow users to back up, restore, or archive the files from their computer. (An archive operation backs up a file, then deletes it from the local disk if the backup is successful.)


NetBackup includes both the server and the client software as follows:

- Server software resides on the computer that manages the storage devices.

- Client software resides on computers that contain data to back up. (Servers also contain client software and can be backed up.)

Figure 1-1 shows an example of a NetBackup storage domain.

Figure 1-1 NetBackup storage domain example



NetBackup accommodates multiple servers that work together under the administrative control of one NetBackup primary server in the following ways:

- The primary server manages backups, archives, and restores. The primary server is responsible for media and device selection for NetBackup. Typically, the primary server contains the NetBackup catalog. The catalog contains the internal databases that contain information about NetBackup backups and configuration.
- Media servers provide additional storage by allowing NetBackup to use the storage devices that are attached to them. Media servers can also increase performance by distributing the network load. Media servers can also be referred to by using the following terms:
 - Device hosts (when tape devices are present)
 - Storage servers (when I/O is directly to disk)

- Data movers (when data is sent to independent, external disk devices like OpenStorage appliances)

During a backup or archive, the client sends backup data across the network to a NetBackup server. The NetBackup server manages the type of storage that is specified in the backup policy.

During a restore, users can browse, then select the files and directories to recover. NetBackup finds the selected files and directories and restores them to the disk on the client.

NetBackup documentation

For a complete list of NetBackup technical documents for each supported release, see the *NetBackup Documentation Landing Page* at the following URL:

<https://www.veritas.com/docs/DOC5332>

The documents are in Adobe® Portable Document Format (PDF), viewable with the Adobe Acrobat Reader. Download the reader from <http://www.adobe.com>.

No responsibility is assumed for the installation and use of the Adobe Acrobat Reader.

NetBackup administration interfaces

NetBackup can be administered with several interfaces. The best choice depends on personal preference and the systems that are available to the administrator.

Table 1-1 NetBackup administration interfaces

Name of interface	Description
NetBackup web user interface	With the NetBackup web user interface (UI), you can view NetBackup activities and manage NetBackup configuration, from a primary server. To start the NetBackup web UI: ■ Users must have a role that is configured for them in NetBackup RBAC. ■ Open a web browser and go to the following URL: <code>https://primaryserver/webui/login</code>
Character-based, menu interface	Run the <code>tpconfig</code> command to start a character-based, menu interface for device management. Use the <code>tpconfig</code> interface from any terminal (or terminal emulation window) that has a <code>termcap</code> or a <code>terminfo</code> definition.

Table 1-1 NetBackup administration interfaces (*continued*)

Name of interface	Description
Command line	NetBackup commands are available on both Windows and UNIX platforms. Enter NetBackup commands at the system prompt or use the commands in scripts. All NetBackup administrator programs and commands require root or administrator user privileges by default. For complete information on all NetBackup commands, see the NetBackup Commands Reference Guide.

About security certificates for NetBackup hosts

NetBackup uses security certificates for authentication of NetBackup hosts. The NetBackup security certificates conform to the X.509 Public Key Infrastructure (PKI) standard. A primary server acts as the NetBackup Certificate Authority (CA) and issues NetBackup certificates to hosts.

NetBackup provides two types of NetBackup host security certificates: Host ID-based certificates and host name-based certificates. Host ID-based certificates are based on Universally Unique Identifiers (UUID) that are assigned to each NetBackup host. The NetBackup primary server assigns these identifiers to the hosts.

NetBackup uses the certificates that are issued from either a NetBackup Certificate Authority or an external certificate authority for host authentication. If you intend to use external certificates on your primary server, you configure the certificates in a post-installation process. The media servers and the clients that use external certificates can either configure external certificates during the installation or upgrade, or after the installation or upgrade.

More information about the post-installation process is available:
https://www.veritas.com/support/en_US/article.100044300

For information on external CA support in NetBackup and external CA-signed certificates, see the [NetBackup Security and Encryption Guide](#).

About setting up the NetBackup Administration Console on UNIX

NetBackup provides two Java-based administration consoles through which the administrator can manage NetBackup. The consoles can be run on either of the following systems:

- Directly on a supported Java-capable UNIX computer by running
`/usr/opensv/java/jnbSA &`

The `jnbSA` command is described in the [NetBackup Commands Reference Guide](#).

Use the `-r` command to connect to the compatible version of the console for the server that you want to administer.

- On a Windows computer that has the **NetBackup Administration Console** installed.
From the multiple versions of consoles installed, select the version of the console that is compatible with the NetBackup server that you want to administer.

Startup procedures and configuration information is explained in the following topics.

Administering remote servers of different versions

The NetBackup web user interface (web UI) is only available for NetBackup 8.1.2 and later. This interface is available on the primary server and supports the version of NetBackup on that server. You do not need to locate and open a specific version as you do with the NetBackup Administration Console. See the documentation for the [NetBackup web UI](#).

The NetBackup server installation provides multiple versions of the **NetBackup Administration Console** to administer remote servers of different versions. When starting the console, select the version of the console that is compatible with the NetBackup server that you want to administer.

Alternatively, from the command line, you can use the `jnbSA` command with the `-r` option to launch the console. For example, to connect to a 8.0 primary server from a 8.1 UNIX primary server, enter the following command on the 8.1 primary server:

```
./jnbSA -r 8.0
```

Several versions of the interface exist. Use the `-h` option and review the `-r` options to find out which versions are supported.

If no `-r` option is specified, the default is the NetBackup version of the current primary server.


Note: To log on to any **NetBackup Administration Console**, your logon credentials must be authenticated from the connecting primary or media server.

Backward compatibility with triple-dot versions

The **NetBackup Administration Console** is backward-compatible between a patch release (x.x.x.x) and a major (x.x) or minor release (x.x.x) that shares the same first and second digits.

For example, the 8.1.x.y console is compatible with a 8.1 primary server. However, a NetBackup 8.1.x.y console cannot administer a 8.0.x.y primary server. See [Figure 1-2](#) for various examples.

Figure 1-2 Examples of supported and unsupported back-level console configurations



See [“Accessing remote servers”](#) on page 1060.

Logging in to the NetBackup Administration Console

Use the **NetBackup Administration Console** to administer and monitor NetBackup operations.

To log in to the NetBackup Administration Console

- 1 On a Windows host where the **NetBackup Administration Console** is installed, select **Start > Programs > Veritas NetBackup > NetBackup version Administration Console**.

On a UNIX computer, use the `jnbSA` command:

```
/usr/opensv/java/jnbSA &
```

The Login screen is displayed.

Note: If the FIPS mode is enabled while you launch the **NetBackup Administration Console**, it is indicated on the title bar of the Login screen.

- 2 The login screen for the **NetBackup Administration Console** displays a name in the **Host name** field.

The default host name is the last host that you successfully logged in to. The drop-down list contains the names of other hosts that you logged in to.

To log in to a different host, type the name of another host.

If the server you enter is a media server or client, the media server or client must have a security certificate installed.

See [“About security certificates for NetBackup hosts”](#) on page 36.

- 3 Select one of the following login options:

- **User name and password**

In the login screen, type your user name and password. To log in to a Windows server, enter both the domain of the server and the user name as follows:

domain_name\user_name

The *domain_name* specifies the domain of the NetBackup host. If the host is not a member of a domain, the *domain_name* is not required.

Note: If the user account is configured for multi-factor authentication on the target host, you must append the one-time password to the password.

For more information on multi-factor authentication, see the *NetBackup Web UI Administrator's Guide*.

- **Windows Active Directory login credentials**

With this option, you can bypass the authentication that is required using the user name and enable Single Sign-on (SSO).

Users with administrative as well as non-administrative privileges can use SSO. The user with administrative privileges needs to right-click and select the **Run as administrator** option while launching the **NetBackup Administration Console**. Selecting this option enables the user to view the console with administrative privileges.

Note the following about SSO:

- The SSO option is available only when both the NetBackup primary server and the NetBackup client are Windows hosts.
 - After the first successful login using SSO, the **Use Active Directory login credentials** option remains in the enabled state for the next login attempt on the same server for the same client.
 - The **NetBackup Administration Console** on a UNIX primary server does not show the option to use the Active Directory credentials.
 - UNIX hosts can log in to the application server using the user name-based authentication.
- **Single sign-on, Certificates, or Smart Cards through the Web UI**
This option enables you to access the NetBackup web UI for authentication using single sign-on (SSO), certificates, or smart cards.
Review the following information
 - This option is disabled if the single sign-on, certificates, or smart cards configurations are not enabled on the primary server. If these

configurations are not available on the primary server, a message is displayed.

- To authenticate through this option, ensure that the primary server is configured for smart cards, user certificates, or SAML 2 FA single sign-on.
See the [Configure NetBackup for Single Sign-On \(SSO\)](#) topic in the [NetBackup Web UI Administrator's Guide](#).
- This option is not supported on NetBackup clients and media servers.
- This option is not supported in NetBackup Access Control (NBAC) mode.
- This option is available through the standalone remote Java consoles and for the primary server.

Users with one of the following permissions can access the **NetBackup Administration Console** using this login option:

- RBAC administrator
- Entry in the `auth.conf` configuration file

The user authentication process with the **Single sign-on, Certificates, or Smart Cards through the Web UI** option:

- The web browser is automatically launched and the NetBackup web UI login page is displayed.
If the browser is not automatically launched, configure the `BROWSER_BINARY_PATH` option in the `nbj.conf` configuration file to launch a browser.
See the [NetBackup Administrator's Guide, Volume I](#).
- Authenticate on the web UI using the certificate, smart cards, or single-sign-on option if these options are configured.
- Once the authentication is successful using the web UI, you can close the web browser and return to the **NetBackup Administration Console** to continue.

- 4 Click **Login** to log in to the NetBackup application server program on the specified server. The interface program continues to communicate through the server that is specified in the logon screen for the remainder of the current session.

Note: If the FIPS mode is enabled while you logon to the **NetBackup Administration Console**, it is indicated on the title bar of the **NetBackup Administration Console**.

See “[User account settings properties](#)” on page 180.

Notes about using the NetBackup Administration Console

- The **NetBackup Administration Console** is best viewed at a screen resolution of 1280 X 1024 or higher. The minimum supported screen resolution to use the console is 1024 X 768.
- The **NetBackup Administration Console** does not support user-defined characters (UDC) and vendor-defined characters (VDC) because of the implementation of Java’s encoding converters.
- On non-English versions of Windows and UNIX systems, the **NetBackup Administration Console** may display non-US ASCII characters incorrectly. This issue can lead to functional failures.
This issue results from a character encoding mismatch between the NetBackup server and the **NetBackup Administration Console**. For a detailed description of the configuration, refer to the following article:
http://www.veritas.com/support/en_US/article.100005338
- To use the **NetBackup Administration Console** on a Windows computer, the Microsoft Windows UAC (User Access Control) feature must be disabled. See the following link for instructions:
<http://windows.microsoft.com/en-us/windows/tum-user-account-control-on-off#1TC=windows-7>
- If there is more than one NetBackup server, the **NetBackup Administration Console** can be run on more than one server at one time. However, if more than one administrator makes changes to the configuration, the results are unpredictable.

See “[Accessing remote servers](#)” on page 1060.

Using the NetBackup web UI

The **NetBackup web UI** provides an interface for the administrator to manage NetBackup.

Table 1-2 Utilities in the left sidebar of the NetBackup web UI

Item	Description
Dashboard	Displays a quick overview of the information that is important to you.
Activity monitor	Displays NetBackup job information and provides the control over the jobs, services, processes, and drives.

Table 1-2 Utilities in the left sidebar of the NetBackup web UI (*continued*)

Item	Description
Recovery	Administrators can use the Recovery node to perform the following kinds of recovery: ■ Regular recovery - Perform server-directed restores of the assets that are protected by policies. Server-directed restores are currently limited to a subset of policy types. Recovery for a specific workload is performed from the Workloads node. For example, to recover VMware assets go to Workloads > VMware. ■ NetBackup catalog recovery. Recovers a catalog backup in a disaster recovery situation.
Protection	Data protection is achieved through policies or protection plans.
Workloads	Contains the supported workloads for NetBackup and tools to manage the workload environment, asset credentials, and recovery.
Storage	This node contains the utilities for managing the media and devices that NetBackup uses to store backups.
Catalog	Search for backup images and perform various actions, including: verify the backup contents, duplicate a backup image, promote a copy, expire a backup image, and import a backup image.
Detection and reporting	This node contains the following tools: ■ Anomaly detection - Detects anomalies in backup metadata. ■ Malware detection - Finds malware in supported backup images and finds the last good-known image that is malware free. ■ Paused protection - Allows NetBackup or authorized users to pause data protection activities. ■ Usage - Displays the primary servers that are configured for capacity licensing and their respective consumption details. ■ Reporting - Use to compile information to verify, manage, and troubleshoot NetBackup operations.
Credential management	Centrally manages the credentials that NetBackup uses to access systems and the workloads that it protects. You can manage credentials for workloads and for systems, client credentials (for NDMP and disk arrays hosts), and External CMS server configurations.
Hosts	Contains the utilities to manage: ■ Deployment management - The main component of VxUpdate that serves as a client or a host upgrade tool. For more information regarding VxUpdate, see the NetBackup Upgrade Guide. ■ Host properties - Use to customize NetBackup configuration options.

Table 1-2 Utilities in the left sidebar of the NetBackup web UI (*continued*)

Item	Description
Resiliency	Integrates NetBackup and Veritas Resiliency Platform to manage your disaster recovery operations.
Security	This node contains the utilities to manage settings for security and hosts: ■ Access keys - Provides access the NetBackup interfaces through API keys and access codes. ■ Certificates - Use to manage NetBackup certificates and view external certificates. ■ Host mappings - Use to carry out NetBackup host operations, such as adding or removing host mappings, resetting a host, or generating a reissue token. ■ Multi-person authorization - Ensures that a second authorized user approves actions before they are performed. ■ RBAC - Use predefined or custom RBAC roles to provide NetBackup users with access to NetBackup, based on their role in your organization. ■ Security events - Use to view the sign-in details for NetBackup users and the user-initiated changes that are made to NetBackup. For more information about Security events, see the NetBackup Security and Encryption Guide. ■ Tokens - Manage the authorization tokens in your NetBackup environment. ■ User sessions - Manage the settings for NetBackup user sessions, terminate user sessions, and unlock a user.
Other licensed utilities	Additional licensed utilities appear under the main NetBackup nodes.

In the upper right-hand corner of the web UI are the following settings.

Table 1-3 Utilities in the top toolbar of the NetBackup web UI

Item	Description
Ticket alerts	Displays a summary of any ticket alerts that are available for multiperson authorization.
Notifications	Displays the most recent events that have occurred in the NetBackup environment.
Help	This menu contains links to the following: ■ NetBackup Help file ■ The NetBackup APIs ■ Log collection. A tool that helps administrators gather evidence to troubleshoot an issue. For more information, see the NetBackup Logging Reference Guide.

Table 1-3 Utilities in the top toolbar of the NetBackup web UI (*continued*)

Item	Description
Settings	This menu contains the following settings: ■ Email notifications - Send email notifications when job failures occur. ■ Global security - Configure security settings for the NetBackup domain. ■ Smart card authentication - Map a smart card or certificate for user validation. ■ Data collector registration - Collect metadata from NetBackup to monitor, manage, and report on NetBackup domains. ■ License management - Manage licenses for NetBackup. ■ Guided setup - Guides you through the process to configure storage, discover virtualization and cloud servers, add protection plans, and protect workloads. ■ NetBackup catalog recovery - Recovers a catalog backup in a disaster recovery situation.
Profile	When you click the profile icon, you can see the following information: ■ Current user's sign in attempts. ■ Password expiration date. ■ NetBackup version of the server. ■ The Approve access request option, to approve an access request that you submitted. ■ The Configure multifactor authentication option, to configure multifactor authentication in NetBackup. ■ The options Add API key or View my API key details, to add your own API key or view the details of your existing API key. ■ In the Profile settings you can add an email address to receive notifications and select ticket events for which you want to receive notifications. ■ The Sign out button, to sign out of the web UI.

Configuring hosts

- [Chapter 2. Configuring Host Properties](#)
- [Chapter 3. Configuring server groups](#)
- [Chapter 4. Enabling support for NAT clients and NAT servers in NetBackup](#)
- [Chapter 5. Configuring host credentials](#)
- [Chapter 6. Managing media servers](#)

Configuring Host Properties

This chapter includes the following topics:

- [About the NetBackup Host properties](#)
- [Access Control properties](#)
- [Active Directory properties](#)
- [Bandwidth properties](#)
- [Backup pool host properties](#)
- [Busy file settings properties](#)
- [Clean up properties](#)
- [Client name properties](#)
- [Client attributes properties](#)
- [Client settings properties for UNIX clients](#)
- [Client settings properties for Windows clients](#)
- [Cloud Storage properties](#)
- [Credential access properties](#)
- [Data Classification properties](#)
- [Default job priorities properties](#)
- [Distributed application restore mapping properties](#)

- Encryption properties
- Enterprise Vault properties
- Enterprise Vault hosts properties
- Exchange properties
- Exclude list properties
- Fibre Transport properties
- Firewall properties
- General server properties
- Global attributes properties
- Logging properties
- Lotus Notes properties
- Media properties
- Network properties
- Network settings properties
- Port ranges properties
- Preferred network properties
- Properties setting in host properties
- RHV access hosts properties
- Resilient network properties
- Resource limit properties
- Restore failover properties
- Retention periods properties
- Scalable Storage properties
- Servers properties
- SharePoint properties
- SLP settings properties
- Throttle bandwidth properties

- [Timeouts properties](#)
- [Universal settings properties](#)
- [User account settings properties](#)
- [UNIX client properties](#)
- [UNIX Server properties](#)
- [VMware access hosts properties](#)
- [Windows client properties](#)
- [Configuration options not found in the host properties](#)
- [About using commands to change the configuration options on UNIX or Linux clients and servers](#)
- [Configuration options for NetBackup servers](#)
- [Configuration options for NetBackup clients](#)

About the NetBackup Host properties

The configuration options within the **Host properties** let an administrator customize NetBackup to meet specific site preferences and requirements.

To change the properties of another client or server, the NetBackup server that you signed in to must be in the **Servers** list on the other system.

See [“Servers properties”](#) on page 159.

See [“Allow access to another server”](#) on page 1058.

For example, if you logged on to *server_1* and want to change a setting on *client_2*, *client_2* must include *server_1* in its **Servers** list.

Some options cannot be configured by using the **NetBackup web UI**.

See [“Configuration options not found in the host properties”](#) on page 186.

Methods to set the NetBackup configuration options

A NetBackup administrator can use one of the following methods to read or set the default configuration options.

Table 2-1 NetBackup Host properties configuration methods

Method	Description
NetBackup Web UI interface	Most properties are listed in the NetBackup web UI in Hosts > Host properties . Depending on the host you want to configure, select the Primary server , Media server , or Client .
Windows registry	Use the <code>nbgetconfig</code> command to obtain a list of configuration entries, and then use <code>nbsetconfig</code> to change the entries in the registry.
<code>bp.conf</code> file	On UNIX, use the <code>nbgetconfig</code> command to obtain a list of configuration entries in the <code>bp.conf</code> file, and then use <code>nbsetconfig</code> to change the entries. The <code>bp.conf</code> file is found in the following location: <pre>/usr/opensv/netbackup/bp.conf</pre> See “About using commands to change the configuration options on UNIX or Linux clients and servers” on page 186. See “Configuration options for NetBackup clients” on page 329.
Command line	Use the <code>nbgetconfig</code> command or <code>bpgetconfig</code> command to obtain a list of configuration entries. Then use <code>nbsetconfig</code> or <code>bpsetconfig</code> to change the options as needed. These commands update the appropriate configuration files on both Windows (registry) and UNIX (<code>bp.conf</code> file) primary servers and clients. Use the <code>nbemmcmd</code> command to modify some options on hosts. Detailed information on these commands is available in the NetBackup Commands Reference Guide.
<code>vm.conf</code> file	The <code>vm.conf</code> file contains configuration entries for media and device management. See the NetBackup Administrator's Guide, Volume II for more information.
Backup, Archive, and Restore client interface	Administrators can specify configuration options for NetBackup clients. See the NetBackup Backup, Archive, and Restore Getting Started Guide.

Connecting to a host to view the host properties

NetBackup displays properties for NetBackup primary servers, media servers, and clients in the **Host properties**.

Use the following procedure to connect to and to view the host properties of a primary server, a media server, or a client.

To connect to and view host properties of primary server, media server, or client

- 1 In the web UI , expand **Hosts > Host properties**.
- 2 Select primary server, media server, or client.
- 3 If necessary, click **Connect**.
- 4 Depending on the host type, select one of the following:
 - **Edit primary server**
 - **Edit media server**
 - **Edit client**

Changing the host properties on multiple hosts at the same time

You can change the host properties for multiple hosts at one time. This can be done by one of the following procedures:

- [Changing multiple hosts in the Host Properties](#)
- [Changing multiple hosts in the Policies utility](#)

Note: In a clustered environment, host properties must be made on each node of the cluster separately.

Changing multiple hosts in the Host Properties

To change the properties on multiple hosts

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties**.
- 2 Select **Primary Server**, **Media Server**, or **Clients**.
- 3 In the right pane, select a host. Hold down the **Shift** key and select another host.
- 4 With multiple hosts still selected, click **Actions > Properties**.

The properties dialog box displays the names of the selected hosts that will be affected by the subsequent host property changes.
- 5 Make changes as necessary.
- 6 Click **OK**.

Changing multiple hosts in the Policies utility

To change the properties on multiple hosts from **Summary of All Policies**

- 1 In the **NetBackup Administration Console**, navigate to **Policies > Summary of All Policies** in the middle pane.
- 2 Under **Summary of All Policies**, expand **Clients**.
- 3 In the right pane, hold down the **Shift** key to select multiple clients.
- 4 With multiple hosts still selected, right-click and select **Host Properties**.
- 5 Change the client properties in the dialog box.
- 6 Click **OK**.

Exporting host properties

Use the following procedure to export the properties of a host.

To export the properties of a host

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Primary Servers, Media Servers, or Clients**.
- 2 Select a host. If you want to select multiple hosts, hold down the **Shift** key and select another host.
- 3 Click **File > Export**.
- 4 In the **Export** dialog box, name the file, browse to the directory where you want to save it, and click **Save**.

Access Control properties

NetBackup Access Control (NBAC) is the legacy access control method for NetBackup and is no longer being updated. It is recommended that you use role-based access control (RBAC) with the web UI. See the [NetBackup Web UI for Administrator's Guide](#).

Use the **Access Control** host properties in the **NetBackup Administration Console** to configure NetBackup Authentication and Authorization. The properties apply to currently selected primary servers, media servers, and clients.

The **NetBackup Product Authentication and Authorization** property displays, regardless of which tab is selected. It determines whether the local system uses access control and how the system uses it.

The **NetBackup Product Authentication and Authorization** property contains the following options.

Table 2-2 NetBackup Product Authentication and Authorization property options

Option	Description
Required	Specifies that the local system should accept requests only from the remote systems that use NetBackup authentication and authorization. Connections from the remote systems that do not use NetBackup authentication and authorization are rejected. Select Required if maximum security is required.
Prohibited	Specifies that the local system should reject connections from any remote system that uses NetBackup authentication and authorization. Select Prohibited if the network is closed and maximum performance is required.
Automatic	Specifies that the local system should negotiate with the remote system about whether to use NetBackup authentication and authorization. Select Automatic if the network contains mixed versions of NetBackup.

For more information about controlling access to NetBackup, see the [NetBackup Security and Encryption Guide](#).

Authentication Domain tab of the Access Control properties

The **Authentication Domain** tab contains the properties that determine which authentication broker a computer uses. A primary server that uses NetBackup authentication and authorization must have at least one authentication domain entry.

If a media server or client does not define an authentication domain, it uses the authentication domains of its primary server.

The **Authentication Domain** tab on the **Access Control** dialog box contains the following properties.

Table 2-3 Authentication Domain tab properties

Property	Description
Available Brokers	Select a broker, then click Find to list all of the available authentication domains.
Available Authentication Domains list	List of available authentication domains.
Add button	Select the authentication domain(s) that this host can use and click Add .
Selected Authentication Domains list	List of the authentication domains that are selected for the host to use.

Table 2-3 Authentication Domain tab properties (*continued*)

Property	Description
Remove button	Select the authentication domain(s) that you no longer want to use and click Remove .

Authorization Service tab of the Access Control properties

The **Authorization Service** tab refers to the authorization service that the local NetBackup server uses. The **Authorization Service** tab does not appear as a property for clients.

The **Authorization Service** tab contains the following properties, which you can configure for a primary or a media server.

Table 2-4 Authorization Service property options

Option	Description
Host name	Specifies the host name or IP address of the authorization service.
Customize the port number of the authorization service	Specifies a nonstandard port number. Select Customize the port number and enter the port number of the authorization service.

Note: Define a host to perform authorization if you configure this tab for a media server to use access control.

Network Attributes tab of the Access Control properties

The **Network Attributes** tab contains a list of networks that are allowed (or not allowed) to use NetBackup authentication and authorization with the local system.

The **Network Attributes** tab on the **Access Control** dialog box contains the following properties:

Networks

The **Networks** property indicates whether specific networks can or cannot use NetBackup authentication and authorization with the local system. The names on the list are relevant only if the **NetBackup Product Authentication and Authorization** property in the **Access Control** dialog box is set to **Automatic** or **Required**.

It is recommended to set **NetBackup Product Authentication and Authorization** property to **Automatic** on the primary server until the clients are configured for access control. Then, change the **NetBackup Product Authentication and Authorization** property on the primary server to **Required**.

If a media server or client does not define a NetBackup Authentication and Authorization network, it uses the networks of its primary server.

NetBackup Product Authentication and Authorization property

The **NetBackup Product Authentication and Authorization property** in this tab determines whether the selected network uses access control and how the network uses it.

See [“Access Control properties”](#) on page 51.

Add Network dialog box

The **Add Network** dialog box contains the following properties.

Table 2-5 Add Network dialog box properties

Property	Description
Host/Domain	Indicates whether the network to be added is a Host name or a Domain name .
Host Details	Specifies that if the network is a host, one of the following items must be entered: ■ The host name of the remote system. (host.domain.com) ■ The IP address of the remote system. (10.0.0.29)
Domain Details	■ Domain Name/IP Enter a dot that is followed by the Internet domain name of the remote systems. (.domain) or the network of the remote system, followed by a dot. (10.0.0.) ■ If the domain is specified by IP, select one of the following items: ■ Bit count Indicates that the mask is based on bit count. Select from between 1 and 32. For example: Mask 192.168.10.10/16 has the same meaning as subnet mask 192.168.20.20:255:255:0.0 ■ Subnet mask Select to enter a subnet mask in the same format as the IP address.

Active Directory properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then and click **Edit client**. Then click **Windows Client > Active Directory**.

The **Active Directory** properties apply to the backup of currently selected Windows Server client. The **Active Directory** properties determine how the backups that allow Active Directory granular restores are performed.

See [“Creating a policy that allows Active Directory granular restores”](#) on page 836.

The **Active Directory** host properties contain the following settings.

Table 2-6 Active Directory properties

Property	Description
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service snapshot provider	Checks snapshots for data corruption. Applies only to snapshots that the Microsoft Volume Shadow Copy Services (VSS) performs. If corrupt data is found and this option is not selected, the job fails. See “Windows open file backup tab of the Client attributes properties” on page 71.
Continue with backup if consistency check fails	Continues the backup job even if the consistency check fails. It may be preferable for the job to continue, even if the consistency check fails. For example, a backup of the database in its current state may be better than no backup at all. Or, it may be preferable for the backup of a large database to continue if it encounters only a small problem.

Bandwidth properties

Use the **Bandwidth** properties to specify network bandwidth limits for the NetBackup clients of the selected primary server.

Note: The **Bandwidth** properties apply only to IPv4 networks. Use the **Throttle Bandwidth** properties to limit IPv6 networks.

See [“Throttle bandwidth properties”](#) on page 174.

The actual limiting occurs on the client side of the backup connection. The bandwidth limits only restrict bandwidth during backups. By default, the bandwidth is not limited.

The **Bandwidth** properties apply to currently selected primary servers.

To manage the **Bandwidth** entries, select one of the following buttons.

Add	Adds an entry to the bandwidth table for each of the selected clients.
Change	Changes an entry to the bandwidth table for each of the selected clients.
Remove	Removes the selected entry from the bandwidth table.

When a backup starts, NetBackup reads the bandwidth limit configuration as configured in the **Bandwidth** host properties. NetBackup then determines the appropriate bandwidth value and passes it to the client. NetBackup computes the bandwidth for each new job based on the number of jobs that are currently running for the IP range. NetBackup does not include local backups in its calculations.

The NetBackup client software enforces the bandwidth limit. Before a buffer is written to the network, client software calculates the current value for kilobytes per second and adjusts its transfer rate if necessary.

Bandwidth limit usage considerations and restrictions

Some usage restrictions apply to the bandwidth limit settings in the **Bandwidth** dialog box. The following table lists the restrictions and describes the specific behaviors that you may need to consider.

Table 2-7 Bandwidth limit usage considerations and restrictions

Client or operation	Bandwidth limit behavior or restrictions
■ Standard ■ MS-Windows	Bandwidth limit is meant primarily for file system backups using Standard and MS-Windows policies. It is not implemented for most other policy and client types.
Local backups	If a server is also a client and data does not go over the network, bandwidth limits have no effect on local backups.
Setting required bandwidth	Bandwidth limits restrict maximum network usage and do not imply required bandwidth. For example, if you set the bandwidth limit for a client to 500 kilobytes per second, the client can use up to that limit. It does not mean, however, that the client requires 500 kilobytes per second.

Table 2-7 Bandwidth limit usage considerations and restrictions (continued)

Client or operation	Bandwidth limit behavior or restrictions
Distributing the workload of active backups	You cannot use bandwidth limits to distribute the backup workload of active backups by having NetBackup pick the most available network segment. NetBackup does not pick the next client to run based on any configured bandwidth limits.

Add Bandwidth Settings dialog box for Bandwidth properties

The **Add Bandwidth Settings** and the **Change Bandwidth Settings** dialog boxes contain the following properties.

From Host	Specifies the beginning of the IP address range of the clients and networks to which the entry applies. For example: 10.1.1.2
To Host	Specifies the end of the IP address range of the clients and networks to which the entry applies. For example: 10.1.1.9
Bandwidth (KB/Sec)	Specifies the bandwidth limitation in kilobytes per second. A value of 0 disables the limits for an individual client or the range of IP addresses covered by the entry. For example, a value of 200 indicates 200 kilobytes per second.

Backup pool host properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Backup host pools**.

The **Backup host pools** properties apply to the backup of the currently selected primary server. A backup host pool is a group of hosts where NetBackup stages the snapshots of the volumes for the backup process to access them. These hosts can be NetBackup clients, media servers, or a primary server. You can create a backup host pool with different versions of NetBackup hosts.

For the hosts that you add to the backup host pool, their volumes are distributed for backup purposes on the backup hosts. This configuration results in a better backup performance.

Note the following important points:

- In a backup host pool you can either have Linux hosts or Windows hosts only. A pool does not support hosts with both platforms.
- All the hosts in the backup host pool must use the same OS version. This way each host has the same version of NFS for consistent backups.
- For backup hosts with a multi-NIC setup, add the host name that is already used on the NetBackup primary server. Do not add an alias name or any other host names in the backup host pool.

Add a backup host pool

To add a backup host pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Backup host pools**.
- 5 Click **Add**.
- 6 Enter the **Backup host pool name**.
- 7 In the **Enter hostname to add to list** box, type the name and click **Add to list**.
- 8 A pool can either have Linux or Windows hosts. To filter the backup hosts in the list, from the **OS type** list select **Windows** or **Linux**.
- 9 From the list, select the hosts that you want to add to the pool.
- 10 Click **Save**.

Add or remove hosts from a backup host pool

To add or remove hosts from a backup host pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Backup host pools**.
- 5 Locate the pool and click **Actions > Edit**.

- 6 A pool can either have Linux or Windows hosts. To filter the backup hosts in the list, from the **OS type** list select **Windows** or **Linux**.
- 7 Select the hosts that you want to include the pool. Or, deselect the hosts you want to remove from the pool.
- 8 Click **Save**.

Delete a backup host pool

You cannot delete a backup host pool if it is part of policy. You must first select a different pool in the policy.

To add or remove hosts from a backup host pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Backup host pools**.
- 5 Locate the pool and click **Actions > Delete > Delete**.

Busy file settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the UNIX client. If necessary click **Connect**, then click **Edit client**. Click **UNIX client > Busy file settings**.

The **Busy file settings** properties define what occurs when NetBackup encounters a busy file during a backup of a UNIX client.

The **Busy file settings** host properties contain the following settings.

Table 2-8 Busy file settings properties

Property	Description
Working directory	Specifies the path to the busy-files working directory. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. By default, NetBackup creates the <code>busy_files</code> directory in the <code>/usr/opensv/netbackup</code> directory.
Administrator email address	Specifies the recipient of the busy-file notification message when the action is set to Send email. By default, the mail recipient is the administrator. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. By default, <code>BUSY_FILE_NOTIFY_USER</code> is not in any <code>bp.conf</code> file and the mail recipient is <code>root</code> .

Table 2-8 Busy file settings properties (*continued*)

Property	Description
Process busy files	Enables busy files to be processed according to the host property settings. NetBackup follows the Busy file settings if it determines that a file changes during a backup. By default, Process busy files is not enabled and NetBackup does not process the busy files. Additional information about busy file processing is available in the NetBackup Administrator's Guide, Volume II.
File action file list	Specifies the absolute path and file name of the busy file. The metacharacters *, ?, [], [-] can be used for pattern matching of file names or parts of file names.
Add	Adds a new file entry. Enter the file and path directly, or browse to select a file.
Add to All	Adds a new file entry for all of the clients currently selected. Enter the file and path directly, or browse to select a file.
Actions > Delete	Deletes the selected file from the file action list.
Retry count	Specifies the number of times to try the backup. The default retry count is 1.
Busy file action	The following options specify which action to take when busy-file processing is enabled. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. ■ Send email sends a busy file notification message to the user that is specified in Administrator email address. ■ Retry the backup retries the backup on the specified busy file. The Retry count value determines the number of times NetBackup tries a backup. ■ Ignore excludes the busy file from busy file processing. The file is backed up, then a log entry that indicates it was busy appears in the All log entries report. The All logs entries report is also available in the NetBackup web UI.

Activating the Busy file settings in host properties

To activate the settings in the **Busy file settings** host properties, use the following procedure.

To activate Busy file settings

- 1 Copy the `bpend_notify_busy` script:

```
/usr/opensv/netbackup/bin/goodies/bpend_notify_busy
```

to the path:

```
/usr/opensv/netbackup/bin/bpend_notify
```

- 2 Set the file access permissions to allow group and others to run `bpend_notify`.
- 3 Configure a policy with a user backup schedule for the busy file backups.

This policy services the backup requests that the repeat option in the actions file generates. The policy name is significant. By default, NetBackup alphabetically searches (uppercase characters first) for the first available policy with a user backup schedule and an open backup window. For example, a policy name of `AAA_busy_files` is selected ahead of `B_policy`.

Clean up properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Clean up**.

The **Clean up** properties manage the retention of various logs and incomplete jobs. The **Clean up** properties apply to primary servers.

The **Clean up** host properties contain the following settings.

Table 2-9 Clean up properties

Property	Description
Keep true image restoration (TIR) information	Specifies the number of days to keep true image restore information on disk. After the specified number of days, the images are pruned (removed). Applies to all policies for which NetBackup collects true image restore information. The default is one day. When NetBackup performs a true image backup, it stores the following images on the backup media: ■ Backed up files ■ True image restore information NetBackup also stores the true image restore information on disk in the following directories: On Windows: <pre>install_path\NetBackup\db\images</pre> On UNIX: <pre>/usr/opensv/netbackup/db/images</pre> NetBackup retains the information for the number of days that this property specifies. Keeping the information on disk speeds up restores. If a user requests a true image restore after the information was deleted from disk, NetBackup retrieves the required information from the media. The only noticeable difference to the user is a slight increase in total restore time. NetBackup deletes the additional information from disk again after one day.
Move restore job from incomplete state to done state	Indicates the number of days that a failed restore job can remain in an Incomplete state. After that time, the Activity monitor shows the job as Done. The default is 7 days. The maximum setting is 365 days. If Checkpoint Restart for restores is used, the Restore retries property allows a failed restore job to be retried automatically. See "Universal settings properties" on page 178. See "Checkpoint restart for restore jobs" on page 677.

Table 2-9 Clean up properties (*continued*)

Property	Description
Move backup job from incomplete state to done state	Indicates the maximum number of hours that a failed backup job can remain in an incomplete state. After that time, the Activity Monitor shows the job as Done. The minimum setting is 1 hour. The maximum setting is 72 hours. The default is 3 hours. When an active job has an error, the job goes into an Incomplete state. In the Incomplete state, the administrator can correct the condition that caused the error. If an Incomplete job does not complete successfully and is moved to the Done state, the job retains the error status. Note: A resumed job reuses the same job ID, but a restarted job receives a new job ID. The job details indicate that the job was resumed or restarted. Note: This property does not apply to suspended jobs. Suspended jobs must be resumed manually before the retention period of the job is met and the image expires. If a suspended job is resumed after the retention period is met, the job fails and is moved to the Done state.
Image cleanup interval	Specifies the maximum interval that can elapse before an image cleanup is run. Image cleanup is run after every successful backup session (that is, a session in which at least one backup runs successfully). If a backup session exceeds this maximum interval, an image cleanup is initiated.
Catalog cleanup wait time	Specifies the minimum interval that can elapse before an image cleanup is run. Image cleanup is not run after a successful backup session until this minimum interval has elapsed since the previous image cleanup.

Client name properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the client. If necessary click **Connect**, then click **Edit client**. Click **Client name**.

The **Client name** property specifies the NetBackup client name for the selected client. The name must match the name the policy uses to back up the client. The only exception is for a redirected restore, where the name must match that of the client whose files are to be restored. The client name is initially set during installation.

The name that is entered here must also match the client name in the **Client attributes** for the primary server. If it does not match, the client cannot browse for its own backups.

Note: Using an IPv6 address as a client name in a policy can cause backups to fail. Specify a host name instead of an IPv6 address.

See “[Client attributes properties](#)” on page 64.

If the value is not specified, NetBackup uses the name that is set in the following locations:

- For a Windows client
In the Network application from the Control Panel.
- For a UNIX client
The name that is set by using the `hostname` command.
The name can also be added to a `$HOME/bp.conf` file on a UNIX client. However, the name is normally added in this manner only for redirected restores. The value in the `$HOME/bp.conf` file takes precedence if it exists.

Client attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**.

The **Client attributes** properties apply to the clients of currently selected primary server.

The **Global client attributes** property applies to all clients, unless overridden as described in the following table.

Table 2-10 Global client attributes

Attribute	Description
Allow client browse	Allows all clients to browse files for restoring. This attribute is overridden if the Browse and restore ability option on the General tab is set to Deny both for a particular clients.
Allow client restore	Allows all clients to restore files. This attribute is overridden if the Browse and restore ability option on the General tab is set to Allow browse only or Deny both .

Table 2-10 Global client attributes (*continued*)

Attribute	Description
Clients	Specifies the list of clients in the client database on the currently selected primary server. A client must be in the client database before you can change the client properties in Client attributes. The client database consists of directories and files in the following directories: Windows: <code>install_path\NetBackup\db\client</code> UNIX: <code>/usr/opensv/netbackup/db/client</code> If a client is not listed in the Clients list, click Add to add a client to the client database. Enter a client name in the text box or select a client. Then click Add. The name that is entered here must match the Client name property for the specific client. If it does not match, the client cannot browse its own backups. See “Client name properties” on page 63. Use the <code>bpclient</code> command to add clients to the client database if dynamic addressing (DHCP) is in use. Additional information about busy file processing is available in the NetBackup Administrator's Guide, Volume II. On UNIX: You also can create, update, list, and delete client entries by using the <code>bpclient</code> command that is located in the following directory: <code>/usr/opensv/netbackup/bin/admincmd</code>
General tab	Specifies how to configure the selected Windows primary servers (clients). See “General tab of the Client attributes properties” on page 66.
Connect options tab	Specifies how to configure the connection between a NetBackup server and a NetBackup client. See “Connect options tab of the Client attributes properties” on page 70.
Windows open file backup tab	Specifies whether a client uses Windows Open File Backup. Also, specifies whether Volume Snapshot Provider or Volume Shadow Copy Service is used as the snapshot provider. See “Windows open file backup tab of the Client attributes properties” on page 71.

General tab of the Client attributes properties

To access this tab, in the web UI select **Hosts > Host properties**. Select the Windows primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**. Then click the **General** tab.

The properties on the **General** tab apply to selected Windows primary servers. The tab appears on the **Client attributes** page.

The **General** tab contains the following properties.

Table 2-11 **General** tab properties

Property	Description
Disable backups until:	Makes the specified clients in the General tab unavailable for backups until the specified date and time. By default, clients are online and included in the policies in which they are listed. When Disable backups until is selected for a client, no jobs are scheduled for that client. Since the client is not part of any job, no backup status is listed for the client. If a client is taken offline, any job is allowed to complete that includes the client and is already running. If a backup or restore job is manually submitted for a client that is offline, the Activity monitor displays the job as failed with a status code 1000 (Client is offline). Note: Changes to this property do not appear in the audit report. The ability to take clients offline is useful in a number of situations. See “Offline option usage considerations and restrictions” on page 68.
Disable restores until:	Makes the specified clients in the General tab unavailable for restores until the specified date and time. By default, clients are online and available for restore.

Table 2-11 **General tab properties** (*continued*)

Property	Description
Maximum data streams	Specifies the maximum number of jobs that are allowed at one time for each selected client. (This value applies to the number of jobs on the client, even if multistreaming is not used.) To change the setting, select Maximum data streams. Then scroll to or enter a value up to 99. The Maximum data streams property interacts with Maximum jobs per client and Limit jobs per policy as follows: ■ If the Maximum data streams property is not set, the limit is either the one indicated by the Maximum jobs per client property or the Limit jobs per policy property, whichever is lower. ■ If the Maximum data streams property is set, NetBackup uses either Maximum jobs per client property. NetBackup uses either Maximum data streams or Limit jobs per policy, whichever is lower. See “Global attributes properties” on page 109. See “Limit jobs per policy (policy attribute)” on page 678.
Browse and restore	Specifies the client permissions to list and restore backups and archives. Select the clients in the General tab of the Client attributes and choose a Browse and restore property. To use the Global client attributes settings, select Use global settings. ■ To allow users on the selected clients to both browse and restore, select Allow both. ■ To allow users on the selected clients to browse but not restore, select Allow browse only. ■ To prevent users on the selected clients from the ability to browse or restore, select Deny both.
Browse and restore scheduled backups	Specifies whether the clients can list and restore from scheduled backups. (This setting does not affect user backups and archives.) This property applies to the privileges that are allowed to a non-Windows administrator or non-root user who is logged into the client. This property also applies to the users that do not have backup and restore privileges. Windows administrators and root users can list and restore from scheduled backups as well as user backups regardless of the Browse and restore scheduled backups setting.

Table 2-11 General tab properties (*continued*)

Property	Description
Deduplication	Specifies the deduplication action for clients if you use the NetBackup Data Protection Optimization Option. For a description of the client-side deduplication options and their actions: See “Where deduplication should occur” on page 69.

Offline option usage considerations and restrictions

The ability to take clients offline is useful in a number of situations. For example, in the event of planned outages or maintenance, client systems can be taken offline to avoid the unnecessary errors that administrators would then need to investigate. This option can also be used to anticipate new clients in the system. You can add them to policies but configure them as offline until they are in place and ready to use.

The following actions can be performed if a client is offline.

Table 2-12 Offline option actions

Type of job or operation	Action or restriction
A client is offline and the job is already in progress.	Offline clients continue to be included in any job.
A client is offline and job retries were started before the client was taken offline.	Job retries continue as normal.
Any duplication job that is associated with a storage lifecycle policy and an offline client.	Continues to run until complete.
Restore jobs	Can be run for offline clients.
The user attempts a manual backup for an offline client.	The backup fails with a status code 1000 (Client is offline). The user can either wait until the client is brought online again or bring the client online manually. Use either the NetBackup web UI or the <code>bpcclient</code> command to do so before resubmitting the manual job.
Archive backups	Not allowed for offline clients.
Administrators restarting or resuming jobs.	Not allowed for offline clients.

Caution: If the primary server is offline, hot catalog backups cannot run.

Where deduplication should occur

The **Deduplication** property specifies the deduplication action for clients if you use the NetBackup Data Protection Optimization Option. More information is available on the client-side deduplication options.

See [Table 2-13](#) on page 69.

The primary server and the clients (that deduplicate their own data) must use the same name to resolve the storage server. The name must be the host name under which the NetBackup Deduplication Engine credentials were created. If they do not use the same name, backups fail. In some environments, careful configuration may be required to ensure that the client and the primary server use the same name for the storage server. Such environments include those that use VLAN tagging and those that use multi-homed hosts.

NetBackup does not support the following for client-side deduplication:

- Multiple copies per each job configured in a NetBackup backup policy. For the jobs that specify multiple copies, the backup images are sent to the storage server and may be deduplicated there.
- NDMP hosts. The backup jobs fail if you try to use client-side deduplication for NDMP hosts.

Table 2-13 Client-side deduplication options

Option	Description
Always use the media server (the default)	Always deduplicates the data on the media server. The default. Jobs fail if one of the following is true: ■ The deduplication services on the storage server are inactive. ■ The deduplication pool is down.
Prefer to use client-side deduplication	Deduplicates the data on the client and then sends it directly to the storage server. NetBackup first determines if the storage server is active. If it is active, the client deduplicates the backup data and sends it to the storage server to be written to disk. If it is not active, the client sends the backup data to a media server, which deduplicates the data.
Always use client-side deduplication	Always deduplicates the backup data on the client and then sends it directly to the storage server. If a job fails, NetBackup does not retry the job.

You can override the **Prefer to use client-side deduplication** or **Always use client-side deduplication** host property in the backup policies.

See “[Client-side deduplication \(policy attribute\)](#)” on page 701.

More information about client deduplication is available in the [NetBackup Deduplication Guide](#).

Connect options tab of the Client attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**. Then click the **Connect options** tab.

The properties in the **Connect options** tab describe how a NetBackup server connects to NetBackup clients. The tab appears on the **Client attributes** page.

The **Connect options** tab contains the following options.

Table 2-14 Connect options tab properties

Property	Description
BPCD connect back	Specifies how daemons are to connect back to the NetBackup Client daemon (BPCD) and contains the following options: ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See “Firewall properties” on page 104. ■ Random port NetBackup randomly chooses a free port in the allowed range to perform the legacy connect-back method. ■ VNETD port NetBackup uses the <code>vnetd</code> port number for the connect-back method.
Ports	Specifies the method that the selected clients should use to connect to the server and contains the following options: ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See “Firewall properties” on page 104. ■ Reserved ports Uses a reserved port number. ■ Non-reserved ports Uses a non-reserved port number.

Windows open file backup tab of the Client attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**. Then click the **Windows open file backup** tab.

Use the settings in this tab only if you want to change the default settings.

By default, NetBackup uses Windows open file backups for all Windows clients. (No clients are listed in the **Client attributes** page.) The server uses the following default settings for all Windows clients:

- Windows open file backup is enabled on the client.
- Microsoft Volume Shadow Copy Service (VSS).
- Snapshots are taken of individual drives (**Individual drive snapshot**) as opposed to all drives at once (**Global drive snapshot**).
- Upon error, the snapshot is terminated (**Abort backup on error**).

Snapshots are a point-in-time view of a source volume. NetBackup uses snapshots to access busy or active files during a backup job. Without a snapshot provider, active files are not accessible for backup.

Table 2-15 Windows open file backup tab properties

Property	Description
Add	Adds a NetBackup client to the list, if you want to change the default settings for Windows open file backups.
Delete	Deletes a client from the list.
Enable Windows open file backup for the selected client	Specifies that Windows open file backup is used for the selected clients. This option functions independently from the Perform Snapshot backups policy option that is available when the Snapshot Client is licensed. If a client is included in a policy that has the Perform Snapshot backups policy option disabled and you do not want snapshots, the Enable Windows open file backups for this client property must be disabled as well for the client. If both options are not disabled, a snapshot is created, though that may not be the intention of the administrator.

Table 2-15 Windows open file backup tab properties (*continued*)

Property	Description
Snapshot Provider	Selects the snapshot provider for the selected clients: ■ Use Veritas Volume Snapshot Provider (VSP) This option is used for back-level versions of NetBackup only. Support for those client versions has ended. ■ Use Microsoft Volume Shadow Copy Service (VSS) Uses VSS to create volume snapshots of volumes and logical drives for the selected clients. For information about how to do Active Directory granular restores when using VSS, see the following topic: See “Active Directory properties” on page 55.
Snapshot usage	Note: The Individual drive snapshot property and the Global drive snapshot property only apply to the non-multistreamed backups that use Windows open file backup. All multistreamed backup jobs share the same volumes snapshots for the volumes in the multistreamed policy. The volume snapshots are taken in a global fashion. Selects how snapshots are made for the selected clients: ■ Individual drive snapshot Specifies that the snapshot should be of an individual drive (default). When this property is enabled, snapshot creation and file backup are done sequentially on a per volume basis. For example, assume that drives C and D are backed up. If the Individual drive snapshot property is selected, NetBackup takes a snapshot of drive C, backs it up, and discards the snapshot. NetBackup then takes a snapshot of drive D, backs it up, and discards the snapshot. Volume snapshots are enabled on only one drive at a time, depending on which drive is to be backed up. This mode is useful when relationships do not have to be maintained between files on the different drives. ■ Global drive snapshot Specifies that the snapshot is of a global drive. All the volumes that require snapshots for the backup job (or stream group for multistreamed backups) are taken at one time. If snapshot creation is not successful, use the Individual drive snapshot option. For example, assume that drives C and D are to be backed up. In this situation, NetBackup takes a snapshot of C and D. Then NetBackup backs up C and backs up D. NetBackup then discards the C and D snapshots. This property maintains file consistency between files in different volumes. The backup uses the same snapshot that is taken at a point in time for all volumes in the backup.

Table 2-15 Windows open file backup tab properties (*continued*)

Property	Description
Snapshot error control	Determines the action to take if there is a snapshot error: Abort backup on error Stops the backup if there is an error during the backup job (after the snapshot is created). The most common reason for a problem after the snapshot is created and is in use by a backup, is that the cache storage is full. If the Abort backup on error property is selected (default), the backup job cancels with a snapshot error status if the backup detects a snapshot issue. This property does not apply to successful snapshot creation. The backup job continues regardless of whether a snapshot was successfully created for the backup job. Disable snapshot and continue Destroys the volume snapshots if the snapshot becomes invalid during a backup. The backup continues with Windows open file backups disabled. Regarding the file that had a problem during a backup—it may be that the file was not backed up by the backup job. The file may not be able to be restored. Note: Volume snapshots typically become invalid during the course of a backup because insufficient cache storage was allocated for the volume snapshot. Reconfigure the cache storage configuration of the Windows open file backup snapshot provider to a configuration that best suits your client's installation.

Client settings properties for UNIX clients

To access this setting, in the web UI select **Hosts > Host properties**. Select the UNIX client. If necessary click **Connect**, then click **Edit client**. Click **UNIX client > Client settings**.

The UNIX **Client settings** properties apply to currently selected NetBackup client running on the UNIX platform.

The UNIX **Client settings** host properties contain the following settings.

Table 2-16 UNIX Client settings properties

Property	Description
Locked file action	Determines what happens when NetBackup tries to back up a file with mandatory file locking enabled in its file mode. Select one of the following options: ■ Wait By default, NetBackup waits for files to become unlocked. If the wait exceeds the Client read timeout host property that is configured on the primary server, the backup fails with a status 41. See “Timeouts properties” on page 175. ■ Skip NetBackup skips the files that currently have mandatory locking set by another process. A message is logged if it was necessary to skip a file.
File compression memory	Specifies the amount of memory available on the client when files are compressed during backup. If you select compression, the client software uses this value to determine how much space to request for the compression tables. The more memory that is available to compress code, the greater the compression and the greater the percentage of computer resources that are used. If other processes also need memory, use a maximum value of half the actual physical memory on a computer to avoid excessive swapping. The default is 0. This default is reasonable; change it only if problems are encountered.
Reset file access time to the value before backup	Specifies that the access time (<i>atime</i>) for a file displays the backup time. By default, NetBackup preserves the access time by resetting it to the value it had before the backup. Note: This setting affects the software and the administration scripts that examine a file's access time. Note: If NetBackup Accelerator is used to perform the backup, this setting is ignored. Accelerator does not record and reset the <i>atime</i> for the files that it backs up. See “Accelerator notes and requirements” on page 707.
Keep status of user-directed backups, archives, and restores	Specifies the number of days to keep progress reports before the reports are deleted. The default is 3 days. The minimum is 0. The maximum is 9,999 days. Logs for user-directed operations are stored on the client system in the following directory: <pre>install_path\NetBackup\logs\user_ops\loginID\logs</pre>

Table 2-16 UNIX Client settings properties (*continued*)

Property	Description
Use VxFS File Change Log (FCL) for incremental backups	Determines if NetBackup uses the File Change Log on VxFS clients. The default is off. See “VxFS file change log (FCL) for incremental backups property” on page 75.
Default cache device path for snapshots	This setting identifies a raw partition available to the copy-on-write process. This raw partition is used when either nbu_snap or VxFS_Snapshot are selected as the snapshot method. The partition must exist on all the clients that are included in the policy.
Add	Adds the file endings to the list of file endings that you do not want to compress. Click Add, then type the file extension. Click Add to add the ending to the list.
Do not compress files ending with these file extensions	Specifies a list of file extensions. During a backup, NetBackup does not compress files with these extensions because the file may already be in a compressed format. Do not use wildcards to specify these extensions. For example, <code>.A1</code> is allowed, but not <code>.A*</code> or <code>.A[1-9]</code> Files that are already compressed become slightly larger if compressed again. If compressed files with a unique file extension already exist on a UNIX client, exclude it from compression by adding it to this list. Corresponds to adding a <code>COMPRESS_SUFFIX =.suffix</code> option to the <code>bp.conf</code> file.

VxFS file change log (FCL) for incremental backups property

The **Use VxFS File Change Log (FCL) for incremental backups** property is supported on all platforms and versions where VxFS file systems support FCL.

The following VxFS file systems support FCL:

- Solaris SPARC platform running VxFS 4.1 or later.
- AIX running VxFS 5.0 or later.
- HP 11.23 running VxFS 5.0 or later.
- Linux running VxFS 4.1 or later.

The File Change Log (FCL) tracks changes to files and directories in a file system. Changes can include the following: files created; links and unlinks; files renamed; data that is appended; data that is overwritten; data that is truncated; extended attribute modifications; holes punched; and file property updates.

NetBackup can use the FCL to determine which files to select for incremental backups, which can potentially save unnecessary file system processing time. The FCL information that is stored on each client includes the backup type, the FCL offset, and the timestamp for each backup.

The advantages of this property depend largely on the number of file system changes relative to the file system size. The performance effect of incremental backups ranges from many times faster or slower, depending on file system size and use patterns.

For example, enable this property for a client on a very large file system that experiences relatively few changes. The incremental backups may complete sooner for the client since the policy needs to read only the FCL to determine what needs to be backed up on the client.

If a file experiences many changes or multiple changes to many files, the time saving benefit may not be as great.

See [“Backup Selections tab”](#) on page 782.

The following items must be in place for the **Use VxFS File Change Log (FCL) for incremental backups** property to work:

- Enable the **Use VxFS File Change Log (FCL) for incremental backups** property for every client that wants NetBackup to take advantage of the FCL.
- Enable the FCL on the VxFS client.
See the Arctera Storage Foundation documentation for information about how to enable the FCL on the VxFS client.
- Enable the **Use VxFS File Change Log (FCL) for incremental backups** property on the clients in time for the first full backup. Subsequent incremental backups need this full backup to stay synchronized.
- Specify the VxFS mount point in the policy backup selections list in one of the following ways:
 - Specify ALL_LOCAL_DRIVES.
 - Specifying the actual VxFS mount point.
 - Specifying a directory at a higher level than the VxFS mount point, provided that option **Cross mount points** is enabled.
See [“Cross mount points \(policy attribute\)”](#) on page 685.

If the policy has **Collect true image restore information** or **Collect true image restore information with move detection** enabled, it ignores the **Use VxFS File Change Log (FCL) for incremental backups** property on the client.

The following table describes the additional options that are available on the VxFS file change log feature.

Table 2-17 VxFS file change log feature options

Option	Description
Activity monitor messages	Displays any messages that note when the file change log is used during a backup as follows: Using VxFS File Change Log for backup of <i>pathname</i> Also notes when full and incremental backups are not synchronized.
Keeping the data files synchronized with the FCL	The data files must be in sync with the FCL for this property to work. To keep the data files synchronized with the FCL on the VxFS client do not turn off the FCL and turn it on again. Note: If NetBackup encounters any errors as it processes the FCL, it switches to the normal files system scan. If this switch occurs, it appears in the Activity monitor.
VxFS administration	Additional VxFS commands are available to administrate the FCL in the Arctera Storage Foundation documentation.

Client settings properties for Windows clients

To access these settings, in the web UI select **Hosts > Host properties**. Select the Windows client and click **Edit client**. Then click **Windows client > Client settings**.

The Windows **Client settings** properties apply to the currently selected Windows client .

The **Windows clients > Client settings** host properties contain the following settings.

Table 2-18 Client settings properties for Windows clients

Property	Description
General level	Enables logs for <code>bpnetd</code>, <code>bpbkar</code>, <code>tar</code>, and <code>nbwin</code>. The higher the level, the more information is written. The default is Minimum logging.

Table 2-18 Client settings properties for Windows clients (*continued*)

Property	Description
Wait time before clearing archive bit	Specifies how long the client waits before the archive bits for a differential incremental backup are cleared. The minimum allowable value is 300 (default). The client waits for acknowledgment from the server that the backup was successful. If the server does not reply within this time period, the archive bits are not cleared. This option applies only to differential-incremental backups. Cumulative-incremental backups do not clear the archive bit.
Use Windows change journal	Note: The Use Windows Change Journal option applies to Windows clients only. This option works together with the Use Accelerator policy attribute and the Accelerator forced rescan schedule attribute. See “Accelerator and the Windows change journal” on page 706. See “Use Accelerator (policy attribute)” on page 703. See “Accelerator forced rescan option (schedule attribute)” on page 742.
Time overlap	Specifies the number of minutes to add to the date range for incremental backups when you use date-based backups. This value compensates for differences in the speed of the clock between the NetBackup client and server. The default is 60 minutes. This value is used during incremental backups when you use the archive bit and when you examine the create time on folders. This comparison is done for archive bit-based backups as well as date-based backups.
Communications buffer size	Specifies the size (in kilobytes) of the TCP and the IP buffers that NetBackup uses to transfer data between the NetBackup server and client. For example, specify 10 for a buffer size of 10 kilobytes. The minimum allowable value is 2, with no maximum allowable value. The default is 128 kilobytes.
User-directed timeouts	Specifies the seconds that are allowed between when a user requests a backup or restore and when the operation begins. The operation fails if it does not begin within this time period. This property has no minimum value or maximum value. The default is 60 seconds.
Perform default search for restore	Instructs NetBackup to search the default range of backup images automatically. The backed up folders and files within the range appear whenever a restore window is opened. Clear the Perform default search for restore check box to disable the initial search. With the property disabled, the NetBackup Restore window does not display any files or folders upon opening. The default is that the option is enabled.

Table 2-18 Client settings properties for Windows clients (*continued*)

Property	Description
TCP level	Enables logs for TCP. Scroll to one of the following available log levels: ■ 0 No extra logging (default) ■ 1 Log basic TCP/IP functions ■ 2 Log all TCP/IP functions ■ 3 Log contents of each read/write Note: Setting the TCP level to 2 or 3 can cause the status reports to be very large. It can also slow a backup or restore operation.
Incrementals	■ Based on timestamp Files that are selected for backup based on the date that the file was last modified. When Use change journal is selected, Based on timestamp is automatically selected. ■ Based on archive bit Note: It is not recommended that you combine differential incremental backups and cumulative incremental backups within the same Windows policy when the incremental backups are based on archive bit. NetBackup include files in an incremental backup only if the archive bit of the file is set. The system sets this bit whenever a file is changed and it normally remains set until NetBackup clears it. A full backup always clears the archive bit. A differential-incremental backup clears the archive bit if the file is successfully backed up. The differential-incremental backup must occur within the number of seconds that the Wait time before clearing archive bit property indicates. A cumulative-incremental or user backup has no effect on the archive bit. If you install or copy files from another computer, the new files retain the date timestamp of the originals. If the original date is before the last backup date on this computer, then the new files are not backed up until the next full backup.
Maximum error messages for single issue	Defines how many times a NetBackup client can send the same error message to a NetBackup server. For example, if the archive bits cannot be reset on a file, this property limits how many times the message appears in the server logs. The default is 10.
Keep status of user-directed backups, archives and restores	Specifies how many days the system keeps progress reports before NetBackup automatically deletes them. The default is 3 days.

How to determine if change journal support is useful in your NetBackup environment

Using NetBackup support for the change journal is beneficial only where the volumes are large and relatively static.

Suitable candidates for enabling NetBackup change journal support are as follows:

- If the NTFS volume contains more than 1,000,000 files and folders and the number of changed objects between incremental backups is small (less than 100,000), the volume is a good candidate for enabling NetBackup change journal support.

Unsuitable candidates for enabling NetBackup change journal support are as follows:

- Support for the change journal is intended to reduce scan times for incremental backups by using the information that is gathered from the change journal on a volume. Therefore, to enable NetBackup change journal support is not recommended if the file system on the volume contains relatively few files and folders. (For example, hundreds of thousands of files and folders.) The normal file system scan is suitable under such conditions.
- If the total number of changes on a volume exceeds from 10% to 20% of the total objects, the volume is not a good candidate for enabling NetBackup change journal support.
- Be aware that virus scanning software can interfere with the use of the change journal. Some real-time virus scanners intercept a file open for read, scan for viruses, then reset the access time. This results in the creation of a change journal entry for every scanned file.

Guidelines for enabling NetBackup change journal support

The following items are guidelines to consider for enabling NetBackup change journal support:

- Change journal support is not offered for user-directed backups. The USN stamps for full and incremental backups in the permanent record do not change.
- NetBackup support for change journal works with checkpoint restart for restores. See [“Checkpoint restart for restore jobs”](#) on page 677.
- Support for change journal is not offered with several NetBackup options.

If **Use Windows change journal** is enabled, it has no effect while you use the following options or products:

- True image restore (TIR) or True image restore with Move Detection
See [“Collect true image restore information \(policy attribute\) with and without move detection”](#) on page 693.

- Synthetic backups
See [“About synthetic backups”](#) on page 842.
- Bare Metal Restore (BMR)
For more information, see the *NetBackup Bare Metal Restore Administrator's Guide*.

See [“How to determine if change journal support is useful in your NetBackup environment”](#) on page 80.

Cloud Storage properties

Note: To access these properties, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Cloud Storage**.

The NetBackup **Cloud Storage** properties apply to the currently selected primary server.

The hosts that appear in this **Cloud Storage** list are available to select when you configure a storage server. The **Service provider** type of your cloud vendor determines whether a service host is available or required.

NetBackup includes service hosts for some cloud storage providers. You can add a new host to the **Cloud Storage** list if the **Service provider** type allows it. If you add a host, you also can change its properties or delete it from the **Cloud Storage** list. (You cannot change or delete the information that is included with NetBackup.)

If you do not add a service host to this **Cloud Storage** list, you can add one when you configure the storage server. The **Service provider** type of your cloud vendor determines whether a **Service host name** is available or required.

Cloud Storage host properties contain the following properties:

Table 2-19 Cloud Storage

Property	Description
Cloud Storage	The cloud storage that corresponds to the various cloud service providers that NetBackup supports are listed here. To add a cloud storage to the Cloud Storage list, click Add. To change properties of a cloud storage that you added, select it in the Cloud Storage list and click Change. To remove a cloud storage that you added, select it in the Cloud Storage list and click Remove.

Table 2-19 Cloud Storage (*continued*)

Property	Description
Associated cloud storage servers for <host>	The cloud storage servers that correspond to the selected cloud storage are displayed. To change the properties of a cloud storage server, select it in the Associated Storage Servers for list and click Change.

For more information about NetBackup cloud storage, see the [NetBackup Cloud Administrator's Guide](#).

Credential access properties

Note: To access these settings, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Credential access**.

Certain NetBackup hosts that are not named as clients in a policy must be enabled to access NDMP or disk array credentials. Use the **Credential access** properties to enter the names of those NetBackup hosts.

The **Credential access** host properties contain the following settings.

Table 2-20 Credential access host properties

Property	Description
NDMP Clients list	To add an NDMP client to the NDMP clients list, click Add . Enter the names of the NDMP hosts that are not named as clients in a policy.

Table 2-20 Credential access host properties (*continued*)

Property	Description
Disk clients list	To add a disk client to the Disk clients list, click Add. Enter the names of the NetBackup hosts that meet all of the following criteria: ■ The host must be designated in a policy as the Off-host backup host in an alternate client backup. ■ The host that is designated as the off-host backup computer must not be named as a client on the Clients tab in any NetBackup policy. ■ The policy for the off-host backup must be configured to use one of the disk array snapshot methods for the EMC CLARiiON, HP EVA, or IBM disk arrays. Note: The credentials for the disk array or NDMP host are specified in the NetBackup web UI. Click Credential management and then click on the Client credentials tab. Note: Off-host alternate client backup is a feature of NetBackup Snapshot Client, which requires a separate license. The NetBackup for NDMP feature requires the NetBackup for NDMP license.

Data Classification properties

To access these settings, in the web UI click **Hosts > Host properties**. Select the server and click **Edit media server** or **Edit primary server**. Then click **Data classification**.

The **Data classification** properties apply to currently selected primary or media server.

Data classifications must be configured in the **Data classification** host properties before storage lifecycle policies can be configured.

See [“Data classifications \(policy attribute\)”](#) on page 668.

Note: Data classifications cannot be deleted. However, the name, description, and the rank can be changed. The classification ID remains the same.

The **Data classification** page contains the following properties.

Table 2-21 Data classification properties

Property	Description
Rank column	The Rank column displays the rank of the data classifications. The order of the data classifications determines the rank of the classification in relationship to the others in the list. The lowest numbered rank has the highest priority. Use the Up and Down buttons to move the classification up or down in the list. To create a new data classification, click Add. New data classifications are added to bottom of the list.
Name column	The Name column displays the data classification name. While data classifications cannot be deleted, the data classification names can be modified. NetBackup provides the following data classifications by default: ■ Platinum (highest rank by default) ■ Gold (second highest rank by default) ■ Silver (third highest rank by default) ■ Bronze (lowest rank by default)
Description column	In the Description, enter a meaningful description for the data classification. Descriptions can be modified.
Data Classification ID	The Data classification ID is the GUID value that identifies the data classification and is generated when a new data classification is added and the host property is saved. . A data classification ID becomes associated with a backup image by setting the Data classification attribute in the policy. The ID is written into the image header. The storage lifecycle policies use the ID to identify the images that are associated with classification. ID values can exist in image headers indefinitely, so data classifications cannot be deleted. The name, description, and rank can change without changing the identity of the data classification.

Adding a Data Classification

Use the following procedures to create or change a data classification.

To add a data classification

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Click **Data classification**.
- 4 Click **Add**.
- 5 Add the name and description.

- 6 Click **Add**.

Note: Data classifications cannot be deleted.

- 7 To change the priority of a classification, select a row and click **Up** or **Down** options.

Default job priorities properties

To access these settings, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Default job priorities**.

The **Default job priorities** host properties let administrators configure the default job priority for different job types.

The job priority can be set for individual jobs in the following utilities:

- In the **Jobs** tab of the **Activity monitor** for queued or active jobs.
See [“Changing the Job Priority dynamically from the Activity Monitor”](#) on page 1012.
- In the **Catalog** utility for verify, duplicate, and import jobs.
- In the **Reports** utility for a Media Contents report job.
- In the **Backup, Archive, and Restore** client interface for restore jobs.

The **Default job priorities** page contains the following properties.

Table 2-22 Default job priorities properties

Property	Description
Job type	The type of job.

Table 2-22 Default job priorities properties (*continued*)

Property	Description
Job priority	The priority that a job has as it competes with other jobs for backup resources. The value can range from 0 to 99999. The higher the number, the greater the priority of the job. A new priority setting affects all the policies that are created after the host property has been changed. A higher priority does not guarantee that a job receives resources before a job with a lower priority. NetBackup evaluates jobs with a higher priority before those with a lower priority. However, the following factors can cause a job with a lower priority to run before a job with a higher priority: ■ To maximize drive use, a low priority job may run first if it can use a drive that is currently loaded. A job with a higher priority that requires that the drive be unloaded would wait. ■ If a low priority job can join a multiplexed group, it may run first. The job with a higher priority may wait if it is not able to join the multiplexed group. ■ If the NetBackup Resource Broker (<code>nbrb</code>) receives a job request during an evaluation cycle, it does not consider the job until the next cycle, regardless of the job priority.

Understanding the job priority setting

NetBackup uses the **Job priority** setting as a guide. Requests with a higher priority do not always receive resources before a request with a lower priority.

NetBackup evaluates the requests sequentially and sorts them based on the following criteria:

- The request's first priority.
- The request's second priority.
- The birth time (when the Resource Broker receives the request).

The first priority is weighted more heavily than the second priority, and the second priority is weighted more heavily than the birth time.

Because a request with a higher priority is listed in the queue before a request with a lower priority, the request with a higher priority is evaluated first. Even though the chances are greater that the higher priority request receives resources first, it is not always definite.

The following scenarios present situations in which a request with a lower priority may receive resources before a request with a higher priority:

- A higher priority job needs to unload the media in a drive because the retention level (or the media pool) of the loaded media is not what the job requires. A lower priority job can use the media that is already loaded in the drive. To maximize drive utilization, the Resource Broker gives the loaded media and drive pair to the job with the lower priority.
- A higher priority job is not eligible to join an existing multiplexing group but a lower priority job is eligible to join the multiplexing group. To continue spinning the drive at the maximum rate, the lower priority job joins the multiplexing group and runs.
- The Resource Broker receives resource requests for jobs and places the requests in a queue before it processes them. New resource requests are sorted and evaluated every 5 minutes. Some external events (a new resource request or a resource release, for example) can also start an evaluation. If the Resource Broker receives a request of any priority while it processes requests in an evaluation cycle, the request is not evaluated until the next evaluation cycle starts.

Distributed application restore mapping properties

To access these settings, in the web UI click **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Distributed application restore mapping**.

Some applications, such as SharePoint, Exchange, and SQL Server distribute and replicate data across multiple hosts. Or, the configuration includes a cluster where communication occurs across multiple nodes. Use the **Distributed application restore mapping** to provide a mapping of the hosts in the database environment so that NetBackup can successfully restore the databases. See the administrator's guide for the database agent for more details.

For example, for a SharePoint farm assume that the farm has two application servers (App1 and App2), one front-end server (FE1) and one SQL database (SQLDB1). The Distributed application restore mapping for this SharePoint server would be as following follows:

Application host	Component host
App1	SQLDB1
App2	SQLDB1
FE1	SQLDB1

The **Distributed application restore mapping** page contains the following properties.

Table 2-23 Distributed application restore mapping properties

Property	Description
Add	This option adds a component host that is authorized to run restores on a SharePoint, Exchange, or SQL Server application host. For SharePoint, NetBackup catalogs backup images under the front-end server name. To allow NetBackup to restore SQL Server back-end databases to the correct hosts in a farm, provide a list of the SharePoint hosts. For Exchange, any operations that use Granular Recovery Technology (GRT) require that you provide a list of the Exchange virtual and the physical host names. You must also include the off-host client and the granular proxy host. For SQL Server, this configuration is required for restores of a SQL Server cluster or a SQL Server availability group (AG). Note: For VMware backups and restores that protect SharePoint, Exchange, or SQL Server, you only need to add the hosts that browse for backups or perform restores. You must also configure a mapping if you use a Primary VM Identifier other than the VM hostname. See the administrator's guide for the database agent for more details. Note: Use either the client's short name or its fully qualified domain name (FQDN). You do not need to provide both names in the list. For more details, see the following: NetBackup for SharePoint Server Administrator's Guide NetBackup for Exchange Server Administrator's Guide NetBackup for SQL Server Administrator's Guide
Actions > Edit	Edits the application host or component host of the currently selected mapping.
Actions > Delete	Deletes the mapping.

Encryption properties

To access these settings, in the web UI click **Hosts > Host properties**. Select the client. If necessary, click **Connect**, then click **Edit client**. Click **Encryption**.

The **Encryption** properties control encryption on the currently selected client.

More information is available in the [NetBackup Security and Encryption Guide](#).

The **Encryption permissions** property indicates the encryption setting on the selected NetBackup client as determined by the primary server.

Table 2-24 Encryption permissions selections

Property	Description
Not allowed	Specifies that the client does not permit encrypted backups. If the server requests an encrypted backup, the backup job ends due to error.
Allowed	Specifies that the client allows either encrypted or unencrypted backups. Allowed is the default setting for a client that has not been configured for encryption.
Required	Specifies that the client requires encrypted backups. If the server requests an unencrypted backup, the backup job ends due to error.

Choose the encryption properties.

Table 2-25 Encryption properties

Property	Description
Use standard encryption	Pertains to the 128-bit and the 256-bit options of NetBackup Encryption.
Client cipher	The following cipher types are available: AES-256-CFB and AES-128-CFB. AES-128-CFB is the default. More information about the ciphers file is available in the NetBackup Security and Encryption Guide .

Additional encryption methods for Windows clients

In addition to NetBackup client and server data encryption, Microsoft Windows clients also have access to methods of encrypting the data on the original disk.

Each of the following methods has its own costs and benefits. NetBackup supports each method for protecting Microsoft Windows clients.

Encrypting File System

The Encrypting File System (EFS) on Microsoft Windows provides file system-level encryption. EFS is a form of encryption where individual files or directories are encrypted by the file system itself.

The technology enables files to be transparently encrypted to protect confidential data from attackers with physical access to the computer. Users can enable encryption on a per-file, per-directory, or per-drive basis. The Group Policy in a Windows domain environment can also mandate some EFS settings.

No NetBackup settings are involved in protecting these encrypted objects. Any object with an encrypted file system attribute is automatically backed up and restored in its encrypted state.

BitLocker Drive Encryption

BitLocker Drive Encryption is a full disk encryption feature included with Microsoft's Windows desktop and server versions.

Disk encryption is a technology which protects information by converting it into unreadable code that cannot be deciphered easily by unauthorized people. Disk encryption uses disk encryption software or hardware to encrypt every bit of data that goes on a disk or a disk volume.

As with EFS, no NetBackup settings are involved to use BitLocker for encryption. Unlike EFS, the encryption layer is invisible to NetBackup, with the data being automatically decrypted and encrypted by the operating system.

NetBackup does nothing to manage the encryption process and therefore backs up and restores the unencrypted data.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must re-enable BitLocker encryption following the restore.

Off-host backup is not supported with volumes that run Windows BitLocker Drive Encryption.

Enterprise Vault properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **Windows Client > Enterprise Vault**.

The **Enterprise Vault** properties apply to currently selected client .

To perform backups and restores, NetBackup must know the user name and password for the account that is used to log on to the Enterprise Vault Server and to interact with the Enterprise Vault SQL database. The user must set the logon account for every NetBackup client that runs backup and restore operations for Enterprise Vault components.

The **Enterprise Vault** host properties contains the following settings.

Table 2-26 Enterprise Vault properties

Property	Description
User name	Specify the user ID for the account that is used to log on to Enterprise Vault (DOMAIN\user name).
Password	Specify the password for the account.
Consistency check before backup	Select what kind of consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation.
Continue with backup if consistency check fails	Continues the backup job even if the consistency check fails. It may be preferable for the job to continue, even if the consistency check fails. For example, a backup of the database in its current state may be better than no backup at all. Or, it may be preferable for the backup of a large database to continue if it encounters only a small problem.

Enterprise Vault hosts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Enterprise Vault hosts**.

The **Enterprise Vault hosts** properties apply to currently selected primary server.

Special configuration is required to allow NetBackup to restore SQL databases to the correct hosts in an Enterprise Vault farm. In the **Enterprise Vault hosts** primary server properties, specify a source and a destination host. By doing so, you specify a source host that can run restores on the destination host.

The **Enterprise Vault hosts** page contains the following properties.

Table 2-27 Enterprise Vault Hosts properties

Option	Description
Add	Adds the source and the destination hosts within the Enterprise Vault configuration. You must provide the name of the Source host and the name of the Destination host .
Actions > Edit	Changes the source host and the destination host.
Actions > Delete	Deletes the entry.

Exchange properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **Windows client > Exchange**.

The **Exchange** properties apply to the currently selected Windows client . For clustered or replicated environments, configure the same settings for all nodes. If you change the attributes for the virtual server name, only the DAG host server is updated.

For complete information on these options, see the [NetBackup for Exchange Server Administrator's Guide](#).

The **Exchange** host properties contain the following settings.

Table 2-28 Exchange properties

Property	Description
Backup option for log files during full backups	Note: This property only applies to MS-Exchange-Server backup policies. Choose which logs to include with snapshot backups: ■ Back up only uncommitted log files (not recommended for replication environments) ■ Back up all log files (including committed log files)
Exchange granular proxy host	Note: This property applies when you duplicate or browse a backup that uses Granular Recovery Technology (GRT). You can specify a different Windows system to act as a proxy for the source client when you duplicate or browse a backup (with <code>bplist</code>) that uses GRT. Use a proxy if you do not want to affect the source client or if it is not available.
Truncate Exchange log files after successful Instant Recovery backup	Note: This property only applies to MS-Exchange-Server backup policies. Enable this option to delete transaction logs after a successful Instant Recovery backup. By default, transaction logs are not deleted for a full Instant Recovery backup that is snapshot only.
Perform consistency check before backup with Microsoft Volume Shadow Copy Service (VSS)	Disable this option if you do not want to perform a consistency check during a DAG backup. If you select Continue with backup if consistency check fails, NetBackup continues to perform the backup even if the consistency check fails.

Table 2-28 Exchange properties (*continued*)

Property	Description
Exchange credentials	Note the following for this property: ■ This property applies to MS-Exchange-Server and VMware backup policies with Exchange recovery. ■ You must configure this property if you want to use GRT. Provide the credentials for the account for NetBackup Exchange operations. This account must have the necessary permissions to perform Exchange restores. The permissions that are required depend on the Exchange version that you have. The account also needs the right to "Replace a process level token."

About the Exchange credentials in the client host properties

The Exchange credentials in the client host properties indicate the account that has necessary permissions to perform Exchange restores. The permissions that are required depend on the Exchange version that you have.

Note the following:

- Credentials are stored in the Credential Management System (CMS).
- To use GRT, configure the Exchange credentials on all granular clients. Alternatively, you can configure the Exchange credentials only on the granular clients that perform restores. In this case, for the entire domain add "Exchange Servers" to the "View-Only Organization Management" role group. Perform this configuration in the Exchange Administration Center (EAC) or in Active Directory. See the following Microsoft article for more information:
<http://technet.microsoft.com/en-us/library/jj657492>
- The account that you configured for the **Exchange credentials** must also have the right to "Replace a process level token."
- For database restores from VMware backups, the Exchange credentials that you provide must have permissions to restore VM files.
- If you want to restore from a VMware snapshot copy that was created with Replication Director, do the following:
 - Provide the Exchange credentials in the **Domain\user** and **Password** fields.
 - Configure the NetBackup Client Service with an account that has access to the CIFS shares that are created on the NetApp disk array.
- If you specify the minimal NetBackup account for the Exchange credentials in the client host properties, NetBackup can back up only active copies of the Exchange databases. If you select **Passive copy only** in the **Exchange**

database backup source field when you create a policy, any backups fail. The failure occurs because the Microsoft Active Directory Service Interface does not provide a list of database copies for a minimal account.

Exclude list properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **Windows client > Exclude list**.

Use the **Exclude list** host properties to create and to modify the exclude list for a Windows client. An exclude list names the files and directories to be excluded from backups.

For information about creating exclude lists for UNIX clients, see the following topic: See [“About excluding files from automatic backups”](#) on page 825.

If more than one exclude or include list exists for a client, NetBackup uses only the most specific one.

For example, assume that a client has the following exclude list:

- An exclude list for a policy and schedule.
- An exclude list for a policy.
- An exclude list for the entire client. This list does not specify a policy or schedule.

In this example, NetBackup uses the first exclude list (for policy and schedule) because it is the most specific.

Exclude and include lists do not determine if an entire drive is excluded when NetBackup determines if a backup job should start.

Normally, a problem does not occur. However, if a policy uses multistreaming and a drive or a mount point is excluded, that job reports an error status when it completes. To avoid the situation, do not use the policy or the policy and the schedule lists to exclude an entire volume.

The **Exclude list** host properties contain the following settings.

Table 2-29 Exclude list properties

Property	Description
Exclude list	Displays the excluded files and directories and the policies and schedules that they apply to. See “Add an entry to an exclude list” on page 96.

Table 2-29 Exclude list properties (*continued*)


Property	Description
Use case-sensitive exclude list	Indicates that the files and directories to exclude are case-sensitive.
Exceptions to the exclude list	Displays any exceptions to the exclude list and the policies and schedules that they apply to. When the policies in this list run, the files and directories in the Exceptions to the exclude list are backed up. Adding an exception can be useful to exclude all files in a directory except one file. See “Add an exception to the exclude list” on page 96. For example, if the file list of items to back up contains <code>/foo</code>, and the exclude list contains <code>/foo/bar</code>, adding <code>/fum</code> to the exceptions list does not back up the <code>/fum</code> directory. However, adding <code>fum</code> to the exceptions list backs up any occurrences of <code>fum</code> (file or directory) that occur within <code>/foo/bar</code>.

About the Add to Exclude List and Add Exceptions to Exclude List dialog boxes

The **Add to Exclude List** dialog box and the **Add Exceptions to Exclude List** dialog box contain the following fields:

Table 2-30 Add to Exclude dialog box

Field	Description
Policy	The policy name that contains the files and the directories that you want to exclude or make exceptions for. You can also select the policy name from the drop-down menu. To exclude or make exceptions for the backup of specific files or directories from all policies, select All Policies .
Schedule	The schedule name that is associated with the files and the directories that you want to exclude or make exceptions for. You can also select the schedule name from the drop-down menu. To exclude or make exceptions for the backups of specific files or directories from all schedules, select All Schedules .
Files/Directories	Either browse or enter the full path to the files and the directories that you want to exclude or make exceptions for.

Figure 2-1 Add to Exclude List properties

Add an entry to an exclude list

Use the following procedure to add an entry to an exclude list for a policy or all policies. When the policies in the exclude list are run, the files and directories that are specified in the list are not backed up.

To add an entry to the exclude list

- 1 Open the NetBackup web UI.
- 2 On the left click **Hosts > Host properties**.
- 3 Select the client.
- 4 If necessary, click **Connect**. Then click **Edit client**.
- 5 Click **Windows clients > Exclude list**.
- 6 Under the Exclude list, click **Add**.
- 7 By default, the file, directory, or path are excluded from **All policies**. Or, type the name of the policy to exclude the items from a specific policy.
- 8 By default, the file, directory, or path are excluded from **All schedules**. Or, type the name of the schedule to exclude the items from a specific policy schedule.
- 9 Enter the file name, directory, or path that you want to exclude from the backups.
- 10 Click **Add**.

Add an exception to the exclude list

Use the following procedure to add an exception to the exclude list for a policy:

To add an exception to the exclude list

- 1** Open the NetBackup web UI.
- 2** On the left click **Hosts > Host properties**.
- 3** Select the client.
- 4** If necessary, click **Connect**. Then click **Edit client**.
- 5** Click **Windows clients > Exclude list**.
- 6** Expand **Exceptions to the exclude list**. Then click **Add**.
- 7** By default, the file, directory, or path is an exception for **All policies**. Or, type the name of the policy to add an exception for a specific policy.
- 8** By default, the file, directory, or path for **All schedules**. Or, type the name of the schedule to add an exception for a specific policy schedule.
- 9** Enter the file name, directory, or path that you want to exclude from the backups.
- 10** Click **Add**.

Syntax rules for exclude lists

It is recommended that you always specify automounted directories and CD-ROM file systems in the exclude list. Otherwise, if the directories are not mounted at the time of a backup, NetBackup must wait for a timeout.

The following syntax rules apply to exclude lists:

- Only one pattern per line is allowed.
- NetBackup recognizes standard wildcard use.
See [“Wildcard use in NetBackup”](#) on page 1055.
See [“NetBackup naming conventions”](#) on page 1054.
- If all files are excluded in the backup selections list, NetBackup backs up only what is specified by full path names in the include list. Files can be excluded by using / or * or by using both symbols together (/*).
- Spaces are considered legal characters. Do not include extra spaces unless they are part of the file name.
For example, if you want to exclude a file named
`C:\testfile` (with no extra space character at the end)
and your exclude list entry is
`C:\testfile` (with an extra space character at the end)
NetBackup cannot find the file until you delete the extra space from the end of the file name.

- End a file path with `\` to exclude only directories with that path name (for example, `C:\users\test\`). If the pattern does not end in `\` (for example, `C:\users\test`), NetBackup excludes both files and directories with that path name.
- To exclude all files with a given name, regardless of their directory path, enter the name. For example:

`test`

rather than

`C:\test`

This example is equivalent to prefixing the file pattern with

`\`

`\*\`

`\*\*\`

`\*\*\*\`

and so on.

The following syntax rules apply only to UNIX clients:

- Do not use patterns with links in the names. For example, assume `/home` is a link to `/usr/home` and `/home/doc` is in the exclude list. The file is still backed up in this case because the actual directory path, `/usr/home/doc`, does not match the exclude list entry, `/home/doc`.
- Blank lines or lines which begin with a pound sign (`#`) are ignored.

Example of a Windows client exclude list

Assume that an exclude list in the **Exclude list** host properties contains the following entries:

`C:\users\doe\john`

`C:\users\doe\abc\`

`C:\users\*\test`

`C:\*\temp`

`core`

Given the exclude list example, the following files, and directories are excluded from automatic backups:

- The file or directory named `C:\users\doe\john`.
- The directory `C:\users\doe\abc\` (because the exclude entry ends with `\`).
- All files or directories named `test` that are two levels beneath `users` on drive C.

- All files or directories named `temp` that are two levels beneath the root directory on drive C.
- All files or directories named `core` at any level and on any drive.

Example of a UNIX exclude list

In this example of a UNIX exclude list, the list contains the following entries:

```
# this is a comment line
/home/doe/john
/home/doe/abc/
/home/*/test
/*temp
core
```

Given the exclude list example, the following files and directories are excluded from automatic backups:

- The file or directory named `/home/doe/john`.
- The directory `/home/doe/abc` (because the exclude entry ends with `/`).
- All files or directories named `test` that are two levels beneath `home`.
- All files or directories named `temp` that are two levels beneath the root directory.
- All files or directories named `core` at any level.

About creating an include list on a UNIX client

To add a file that is eliminated with the exclude list, create a `/usr/opensv/netbackup/include_list` file. The same syntax rules apply as for the exclude list.

Note: Exclude and include lists do not apply to user backups and archives.

To illustrate the use of an include list, we use the example from the previous discussion. The exclude list in that example causes NetBackup to omit all files or directories named `test` from all directories beneath `/home/*/test`.

In this case, add a file named `/home/jdoe/test` back into the backup by creating an `include_list` file on the client. Add the following to the `include_list` file:

```
# this is a comment line
/home/jdoe/test
```

To create an include list for a specific policy or policy and schedule combination, use a `.policyname` or `.policyname.schedulename` suffix. The following are two examples of include list names for a policy that is named `wkstations` that contains a schedule that is named `fulls`.

```
/usr/opensv/netbackup/include_list.workstations  
/usr/opensv/netbackup/include_list.workstations.fulls
```

The first file affects all scheduled backups in the policy that is named `wkstations`. The second file affects backups only when the schedule is named `fulls`.

For a given backup, NetBackup uses only one include list: the list with the most specific name. Given the following two files:

```
include_list.workstations  
include_list.workstations.fulls
```

NetBackup uses only `include_list.workstations.fulls` as the include list.

Traversing excluded directories

An exclude list can indicate a directory for exclusion, while the client uses an include list to override the exclude list. NetBackup traverses the excluded directories if necessary, to satisfy the client's include list.

Assume the following settings for a Windows client:

- The backup policy backup selection list indicates `ALL_LOCAL_DRIVES`. When a scheduled backup runs, the entire client is backed up.
The entire client is also backed up if the backup selection list consists of only:
/
 - The exclude list on the client consists of only: *An exclude list of * indicates that all files are excluded from the backup.
- However, since the include list on the Windows client includes the following file:
`C:\WINNT`, the excluded directories are traversed to back up `C:\WINNT`.
If the include list did not contain any entry, no directories are traversed.

In another example, assume the following settings for a UNIX client:

- The backup selection list for the client consists of the following: /
- The exclude list for the UNIX client consists of the following: /
- The include list of the UNIX client consists of the following directories:
/data1
/data2
/data3

Because the include list specifies full paths and the exclude list excludes everything, NetBackup replaces the backup selection list with the client's include list.

Fibre Transport properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the host. If necessary, click **Connect**. Then click **Edit primary server**, or **Edit media server**, or **Edit client**. Click **Fibre Transport**.

Note that any usage preferences that you configure for an individual SAN Client in **Storage > SAN Clients** override the Fibre Transport properties that you configure in **Host properties**.

NetBackup Fibre Transport properties control how your Fibre Transport media servers and SAN clients use the Fibre Transport service for backups and restores. The **Fibre Transport** properties apply to the host type that you select, as follows:

Table 2-31 Host types for Fibre Transport properties

Host type	Description
Primary server	Global Fibre Transport properties that apply to all SAN clients.
Media server	The Fibre Transport Maximum concurrent Fibre Transport connections property applies to the Fibre Transport media server that you select.
Client	The Fibre Transport properties apply to the SAN client that you select. The default values for clients are the global property settings of the primary server. Client properties override the global Fibre Transport properties.

The **Fibre Transport** properties contain the following settings. All properties are not available for all hosts. In this table, the Fibre Transport device is an HBA port on a Fibre Transport media server. The port carries the backup and restore traffic. A media server may have more than one Fibre Transport device.

Table 2-32 Fibre Transport properties

Property	Description
Maximum concurrent Fibre Transport connections Maximum concurrent FT connections	This property displays only when you select a Fibre Transport media server . This property specifies the number of Fibre Transport connections to allow to the selected media server or media servers. A connection is equivalent to a job. If no value is set, NetBackup uses the following defaults: ■ For NetBackup Appliance model 5330 and later: 32 ■ For NetBackup Appliance model 5230 and later: 32 ■ For NetBackup Fibre Transport media servers: 8 times the number of fast HBA ports on the media server plus 4 times the number of slow HBA ports. A fast port is 8 GB or faster, and a slow port is less than 8 GB. You can enter up to the following maximum connections for the media server or servers to use: ■ On a Linux Fibre Transport media server host: 40. The recommended setting is 32 or fewer connections concurrently on Linux. On Linux hosts, you can increase that maximum by setting a NetBackup touch file, <code>NUMBER_DATA_BUFFERS_FT</code>. See “About Linux concurrent FT connections” on page 103. ■ For NetBackup Appliance model 5330 and later: 40. ■ For NetBackup Appliance model 5230 and later: 40. ■ On a Solaris Fibre Transport media server host: 64. NetBackup supports 644 buffers per media server for Fibre Transport. To determine the number of buffers that each connection uses, divide 644 by the value you enter. If you have buffers per connection you see better performance for each connection.
Use defaults from the primary server configuration	This property displays only when you select a client . This property specifies that the client follow the properties as they are configured on the primary server.
Preferred	The Preferred property specifies to use an FT device if one is available within the configured wait period in minutes. If an FT device is not available after the wait period elapses, NetBackup uses a LAN connection for the operation. For the global property that is specified on the primary server, the default is Preferred.

Table 2-32 Fibre Transport properties (*continued*)

Property	Description
Always	The Always property specifies that NetBackup should always use an FT device for backups and restores of SAN clients. NetBackup waits until an FT device is available before it begins the operation. However, an FT device must be active and up. If not, NetBackup uses the LAN. An FT device may not exist because none is active, none have been configured, or the SAN Client license expired.
Fail	The Fail property specifies that NetBackup should fail the job if an FT device is not active and up. If the FT devices are active but busy, NetBackup waits until a device is available and assigns the next job to the device. An FT device may not exist because none are active, none have been configured, or the SAN Client license expired.
Never	The Never property specifies that NetBackup should never use a Fibre Transport pipe for backups and restores of SAN clients. NetBackup uses a LAN connection for the backups and restores. If you specify Never for the primary server, Fibre Transport is disabled in the NetBackup environment. If you select Never, you can configure Fibre Transport usage on a per-client basis. If you specify Never for a SAN client, Fibre Transport is disabled for the client.

NetBackup provides one finer level of granularity for Fibre Transport. SAN client usage preferences override the FT properties that you configure through **Host properties**.

For more information about NetBackup Fibre Transport, see the [NetBackup SAN Client and Fibre Transport Guide](#).

About Linux concurrent FT connections

NetBackup uses the **Maximum concurrent FT connections** setting in the **Fibre transport** host property to configure the number of concurrent connections to a Fibre transport media server, up to the total that is allowed per host.

See “[Fibre Transport properties](#)” on page 101.

If the total number of concurrent connections on Linux is too low for your purposes, you can increase the total number of concurrent connections. The consequence is that each client backup or restore job uses fewer buffers, which means that each job is slower because of fewer buffers. To increase the number of concurrent connections, reduce the number of buffers per connection. To do so, create the following file and include one of the supported values from [Table 2-33](#) in the file:

/usr/opensv/netbackup/db/config/NUMBER_DATA_BUFFERS_FT

Table 2-33 shows the values that NetBackup supports for the `NUMBER_DATA_BUFFERS_FT` file. NetBackup supports 644 buffers per media server for Fibre transport.

Table 2-33 Supported values for buffers per FT connection

NUMBER_DATA_BUFFERS_FT	Total concurrent connections: NetBackup 5230 and 5330 and later appliances	Total concurrent connections: Linux FT media server
16	40	40
12	53	53
10	64	64

If you want, you then can limit the number of connections for a media server with the **Maximum concurrent FT connections** setting in the **Fibre transport** host properties.

Firewall properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server or media server. If necessary click **Connect**, then click **Edit primary server** or **Edit media server**. Click **Firewall**.

The **Firewall** properties determine how the selected primary servers and media servers connect to the legacy services that run on that NetBackup host.

Servers are added to the **Hosts** list of the **Firewall** properties. To configure port usage for clients, see the **Client attributes** properties.

See “[Client attributes properties](#)” on page 64.

The **Firewall** host properties contain the following settings.

Table 2-34 Firewall properties

Property	Description
Default connect options	By default, the Default connect options include firewall-friendly connect options including the fewest possible ports to open. The default options can be set differently for an individual server or client with the settings in Attributes for selected hosts. To change the default connect options for the selected server or client, click Edit. These properties correspond to the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option.
Hosts	You can configure different default connect options for the hosts that are displayed in this list. - Click Add to add a host to the Hosts list. You must add a host name to the list before you can configure different settings for that host. Servers do not automatically appear on the list. - To configure different settings for a host, select the host name in the Hosts list. Then select the connect options in the Attributes for selected hosts section. - To remove the host from the list, locate a host name in the list. Then click Delete.
Attributes for selected hosts	This section displays the connect options for the selected server. To change the connection options for a server, first select the host name in the Hosts list. These properties correspond to the <code>CONNECT_OPTIONS</code> configuration option.
BPCD connect back	This property specifies how daemons are to connect back to the NetBackup Client daemon (<code>BPCD</code>) as follows: Use default connect options (An option for individual hosts) Use the methods that are specified under Default connect options. Random port NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method. VNETD port This method requires no connect-back. The Cohesity Network Daemon (<code>vnetd</code>) was designed to enhance firewall efficiency with NetBackup during server-to-server and server-to-client communications. The server initiates all <code>bpcd</code> socket connections. Consider the example in which <code>bpbrm</code> on a media server initially connects with <code>bpcd</code> on a client. The situation does not pose a firewall problem because <code>bpbrm</code> uses the well-known PBX or <code>vnetd</code> port.

Table 2-34 Firewall properties (*continued*)

Property	Description
Ports	Select whether a reserved or non-reserved port number should be used to connect to the host name: ■ Use default connect options (An option for individual hosts) Use the methods that are specified under Default attributes. ■ Reserved ports Connect to the host name by a reserved port number. ■ Non-reserved ports Connect to the host name by a non-reserved port number. To configure port usage for clients, see the Client attributes properties.

General server properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server or media server. If necessary click **Connect**, then click **Edit primary server** or **Edit media server**. Click **General server**.

The **General server** properties apply to the selected primary server or media server.

The **General server** page contains the following properties.

Table 2-35 General server properties

Property	Description
Delay on multiplexed restores	This property specifies how long the server waits for additional restore requests of multiplexed images on the same tape. All of the restore requests that are received within the delay period are included in the same restore operation (one pass of the tape). The default is a delay of 30 seconds.
Check the capacity of disk storage units every	This property applies to the disk storage units of 6.0 media servers only. Subsequent releases use internal methods to monitor disk space more frequently.

Table 2-35 General server properties (*continued*)

Property	Description
Must use local drive	This property appears for primary servers only, but applies to all media servers as well. This property does not apply to NDMP drives. If a client is also a media server or a primary server and Must use local drive is selected, a local drive is used to back up the client. If all drives are down, another can be used. This property increases performance because backups are done locally rather than sent across the network. For example, in a SAN environment a storage unit can be created for each SAN media server. Then, the media server clients may be mixed with other clients in a policy that uses ANY AVAILABLE storage unit. When a backup starts for a client that is a SAN media server, the backups go to the SAN connected drives on that server.
Use direct access recovery for NDMP restores	By default, NetBackup for NDMP is configured to use Direct Access Recovery (DAR) during NDMP restores. DAR can reduce the time it takes to restore files by allowing the NDMP host to position the tape to the exact location of the requested file(s). Only the data that is needed for those files is read. Clear this check box to disable DAR on all NDMP restores. Without DAR, NetBackup reads the entire backup image, even if only a single restore file is needed.
Enable message-level cataloging when duplicating Exchange images that use Granular Recovery Technology	This option performs message-level cataloging when you duplicate Exchange backup images that use Granular Recovery Technology (GRT) from disk to tape. To perform duplication more quickly, you can disable this option. However, then users are not able to browse for individual items on the image that was duplicated to tape. See the NetBackup for Exchange Administrator's Guide.

Table 2-35 General server properties (*continued*)

Property	Description
Media host override list	Specific servers can be specified in this list as servers to perform restores, regardless of where the files were backed up. (Both servers must be in the same primary and media server cluster.) For example, if files were backed up on media server A, a restore request can be forced to use media server B. The following items describe situations in which the capability to specify servers is useful: ■ Two (or more) servers share a robot and each have connected drives. A restore is requested while one of the servers is either temporarily unavailable or is busy doing backups. ■ A media server was removed from the NetBackup configuration, and is no longer available. To add a host to the Media host override list, click Add. To change an entry in the list, select a host name, then click Actions > Edit. Configure the following options: ■ Original backup server Enter the name of the server where the data was backed up originally. ■ Restore server Enter the name of the server that is to process future restore requests.

Forcing restores to use a specific server

Use the following procedure to force restores to use a specific server.

To force restores to use a specific server

- 1 If necessary, physically move the media to the host to answer the restore requests, then update the NetBackup database to reflect the move.
- 2 Modify the NetBackup configuration on the primary server.
 - Open the NetBackup web UI and sign into the primary server.
 - On the left, click **Host > Host properties**.
 - Select the primary server.
 - If necessary, click **Connect**. Then click **Edit primary server**.
 - Click **General server**.

- Add the original backup media server and the restore server to the **Media host override** list.
- 3 Stop and restart the NetBackup Request Daemon (`bprd`) on the primary server.
- This process applies to all storage units on the original backup server. Restores for any storage unit on the **Original backup server** go to the server that is listed as the **Restore server**.
- To revert to the original configuration for future restores, delete the line from the **Media host override** list.

Global attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Global attributes**.

The **Global attributes** properties apply to currently selected primary servers. These properties affect all operations for all policies and clients. The default values are adequate for most installations.

The **Global attributes** page contains the following properties.

Table 2-36 Global attributes properties

Property	Description
Job retry delay	This property specifies how often NetBackup retries a job. The default is 10 minutes. The maximum is 60 minutes; the minimum is 1 minute.

Table 2-36 Global attributes properties (*continued*)

Property	Description
Maximum jobs per second	This property specifies the throttle on the maximum number of backup jobs that are allowed to go from the Queued to Active state per second. By default, the value of this property is 0, which means no throttling occurs. After the maximum number of jobs is reached in one second, subsequent jobs will remain in the Queued state. In the next second, jobs are released in a first-in-first-out order from the Queued state until the maximum jobs value is reached again or until all throttled jobs or new jobs have been made active. This property can be used to smooth out the resource utilization curve. It is particularly useful when backup windows open and a large number of jobs are scheduled to start within a short time period. This value supersedes the <code>DBM_NEW_IMAGE_DELAY</code> configuration value found here: https://www.veritas.com/support/en_US/article.100047119 If <code>DBM_NEW_IMAGE_DELAY</code> is configured and the maximum jobs per second throttle is the default value, the <code>DBM_NEW_IMAGE_DELAY</code> will be converted to an equivalent jobs-per-second value. It will not modify the configuration. For example, if <code>DBM_NEW_IMAGE_DELAY</code> was set to 333ms, the NetBackup Job Manager will use a maximum jobs per second throttle of 3. If the user then configured the maximum jobs throttle to 2 per second, the configured <code>DBM_NEW_IMAGE_DELAY</code> would be ignored. Note: This throttle only affects the number of backup jobs that the NetBackup Job Manager allows to start in one second. It does not affect other job types like restores, archives, duplications, or replications. It does not affect the maximum number of concurrent jobs.
Maximum jobs per client	This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently. The default is one job. NetBackup can process concurrent backup jobs from different policies on the same client only in the following situations: ■ More than one storage unit available ■ One of the available storage units can perform more than one backup at a time. See “About constraints on the number of concurrent jobs” on page 112.
Policy update interval	This property specifies how long NetBackup waits to process a policy after a policy is changed. The interval allows the NetBackup administrator time to make multiple changes to the policy. The default is 10 minutes. The maximum is 1440 minutes; the minimum is 1 minute.
Compress catalog interval	This property specifies how long NetBackup waits after a backup before it compresses the image catalog file.

Table 2-36 Global attributes properties (*continued*)

Property	Description
Schedule backup attempts	NetBackup considers the failure history of a policy to determine whether or not to run a scheduled backup job. The Schedule backup attempts property sets the timeframe for NetBackup to examine. This property determines the following characteristics for each policy: ■ How many preceding hours NetBackup examines to determine whether to allow another backup attempt (retry). By default, NetBackup examines the past 12 hours. ■ How many times a backup can be retried within that timeframe. By default, NetBackup allows two attempts. Attempts include the scheduled backups that start automatically or the scheduled backups that are user-initiated. Consider the following example scenario using the default setting 2 tries every 12 hours: ■ Policy_A runs at 6:00 P.M.; Schedule_1 fails. ■ Policy_A is user-initiated at 8:00 P.M.; Schedule_2 fails. ■ At 11:00 P.M., NetBackup looks at the previous 12 hours. NetBackup sees one attempt at 6:00 P.M. and one attempt at 8:00 P.M. The Schedule backup attempts setting of two has been met so NetBackup does not try again. ■ At 6:30 A.M. the next morning, NetBackup looks at the previous 12 hours. NetBackup sees only one attempt at 8:00 P.M. The Schedule backup attempts setting of two has not been met so NetBackup tries again. If a schedule window is not open at this time, NetBackup waits until a window is open. Note: This attribute does not apply to user backups and archives.
Maximum vault jobs	This property specifies the maximum number of vault jobs that are allowed to be active on the primary server. The greater the maximum number of vault jobs, the more system resources are used. If the active vault jobs limit is reached, subsequent vault jobs are queued and their status is shown as Queued in the Activity monitor. If a duplication job or eject job waits, its status is shown as Active in the Activity monitor.
Administrator email address property	This property specifies the addresses where NetBackup sends notifications of scheduled backups or administrator-directed manual backups. To send the information to more than one administrator, separate multiple email addresses by using a comma, as follows: <pre>useraccount1@company.com,useraccount2@company.com</pre> More information is available on the configuration requirements for email notifications. See “Send notifications to the backup administrator about failed backups” on page 1047.


About constraints on the number of concurrent jobs

Specify any number of concurrent jobs within the following constraints.

Table 2-37 Constraints on concurrent jobs

Constraint	Description
Number of storage devices	NetBackup can perform concurrent backups to separate storage units or to drives within a storage unit. For example, a single Media Manager storage unit supports as many concurrent backups as it has drives. A disk storage unit is a directory on disk, so the maximum number of jobs depends on system capabilities.
Server and client speed	Too many concurrent backups on an individual client interfere with the performance of the client. The best setting depends on the hardware, operating system, and applications that are running. The Maximum jobs per client property applies to all clients in all policies. To accommodate weaker clients (ones that can handle only a small number of jobs concurrently), consider using one of the following approaches: - Set the Maximum data streams property for those weaker clients appropriately. (Open the host properties for the primary server. Then click Client attributes > General tab.) See “General tab of the Client attributes properties” on page 66. - Use the Limit jobs per policy policy setting in a client-specific policy. (A client-specific policy is one in which all clients share this characteristic). See “Limit jobs per policy (policy attribute)” on page 678.
Network loading	The available bandwidth of the network affects how many backups can occur concurrently. The load might be too much for a single Ethernet. For loading problems, consider backups over multiple networks or compression. A special case exists to back up a client that is also a server. Network loading is not a factor because the network is not used. Client and server loading, however, is still a factor.

Figure 2-2 Maximum jobs per client



Note: Catalog backups can run concurrently with other backups. To do so, set the **Maximum jobs per client** value to greater than two for the primary server. The higher setting ensures that the catalog backup can proceed while the regular backup activity occurs.

Setting up mailx email client

NetBackup supports setting up email notifications by using mailx client.

To set up a mailx email client

- 1 Navigate to the `/etc/mail.rc` location.
- 2 Edit the file to add the SMTP server settings.

For example, set

```
smtp=<Your_SMTP_Server_Hostname>:<SMTP_SERVER_PORT>
```

Logging properties

To access the Logging properties, in the web UI select **Hosts > Host properties**. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Logging**.

The logging settings determine the behavior for NetBackup logging on the primary server, media server, and the clients:

- Overall logging level or global logging level for all NetBackup processes.
- Overrides for the specific processes that use legacy logging.
- Logging levels for the services that use unified logging.
- Logging for critical processes.
- On clients, the logging level for database applications.
- Log retention settings for NetBackup and for NetBackup Vault (if it is installed).

All NetBackup processes use either unified logging or legacy logging. You can set a global or a unique logging level for certain processes and services. Retention levels limit the size of the log files or (for the primary server) the number of days the logs are kept. If you use NetBackup Vault, you can select separate logging retention settings for that option.

For complete details on logging, see the [NetBackup Logging Reference Guide](#).

Table 2-38 Logging properties

Property	Description
Global logging level	This setting establishes a global logging level for all processes that are set to Same as global. The Global logging level affects the legacy and unified logging level of all NetBackup processes on the server or client. This setting does not affect the following logging processes: ■ PBX logging See the NetBackup Troubleshooting Guide for more information on how to access the PBX logs. ■ Media and device management logging (vmd, ltid, avrd, robotic daemons, media manager commands)
Process-specific overrides	These settings let you override the logging level for the specific processes that use legacy logging.
Debug logging levels for NetBackup services	These settings let you manage the logging level for the specific services that use unified logging.
Logging for critical processes	The option lets you enable logging for the critical processes: ■ Primary server processes: bprd and bpbm. ■ Media server processes: bpbm, bptm, and bpdm. ■ Client process: bpfis Note the following: ■ If you enable Logging for critical processes, also enable the option Maximum log size. If you disable this option it may adversely affect NetBackup operations. ■ This option sets the log retention to the default log size. ■ Clicking Restore to defaults does not modify the Logging for critical processes or the Maximum log size options. ■ To disable the logging for critical processes, modify the logging levels for those processes.
Retention period	Specifies the length of time NetBackup keeps information from the error catalog, job catalog, and debug logs. Note that NetBackup derives its reports from the error catalog. The logs can consume a large amount of disk space, so do not keep the logs any longer than necessary. The default is 28 days. Note: This setting is not applicable for Cloud Scale.

Table 2-38 Logging properties (*continued*)

Property	Description
Maximum log size	Specifies the size of the NetBackup logs that you want to retain. When the NetBackup log size grows to this value, the older logs are deleted. ■ For primary and media servers, the recommended value is 25 GB or greater. ■ For clients, the recommended value is 5 GB or greater. Note: This setting is not applicable for Cloud Scale.
Vault logs retention period	If NetBackup Vault is installed, select the number of days to keep the Vault session directories, or select Forever .

Logging levels

You can choose to apply the same logging level for all NetBackup processes. Or, you can select logging levels for specific processes or services.

Table 2-39 Logging level descriptions

Logging level	Description
Same as global	The process uses the same logging level as the Global logging level .
No logging	No log is created for the process.
Minimum logging (default)	A small amount of information is logged for the process. Use this setting unless advised otherwise by Cohesity Technical Support. Other settings can cause the logs to accumulate large amounts of information.
Levels 1 through 4	Progressively more information is logged at each level for the process.
5 (Maximum)	The maximum amount of information is logged for the process.

Global logging level

This setting controls the logging level for all processes and for those processes that are set to **Same as global**. You can control the logging level for some NetBackup processes individually.

See [the section called “Overrides for legacy logging levels”](#) on page 116.

See [the section called “Unified logging levels for the primary server”](#) on page 116.

Overrides for legacy logging levels

These logging levels apply to legacy processes logging. The logging levels that are displayed depend on the type of host (primary, media, or client).

Table 2-40 Logging level overrides for legacy processes

Service	Description	Primary server	Media server	Client
BPBRM logging level	The NetBackup backup and restore manager.	X	X	
BPDM logging level	The NetBackup disk manager.	X	X	
BPTM logging level	The NetBackup tape manager.	X	X	
BPJOB logging level	The NetBackup Jobs Database Management daemon. This setting is only available for the primary server.	X		
BPDBM logging level	The NetBackup database manager.	X		
BPRD logging level	The NetBackup Request Daemon.	X		
Database logging level	The logging level for database agent logs. For details on which logs to create and refer to, see the guide for the specific agent.			X

Unified logging levels for the primary server

These logging levels apply to NetBackup services logging and are only available for the primary server.

Table 2-41 Logging levels for NetBackup services

Service	Description
Policy execution manager	The Policy execution manager (NBPPEM) creates policy and client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, NBPPEM is notified and the appropriate policy and client tasks are updated.
Job manager	The Job Manager (NBJM) accepts the jobs that the Policy Execution Manager submits and acquires the necessary resources.
Resource broker	The Resource Broker (NBRB) makes the allocations for storage units, tape drives, client reservations.

Logging values in the registry, bp.conf file, and unified logging

You can also set logging values in the Windows registry, the bp.conf file, or in unified logging.

Table 2-42 Logging levels and their values

Logging level	Legacy logging - Windows registry	Legacy logging - bp.conf	Unified logging
Minimum logging	Hexadecimal value of 0xffffffff.	<code>VERBOSE = 0 (global)</code> <code>processname_VERBOSE = 0</code> If the global VERBOSE value is set to a value other than 0, an individual process can be decreased by using the value -1. For example, <code>processname_VERBOSE = -1.</code>	1
No logging	Hexadecimal value of 0xffffffffe.	<code>VERBOSE=-2 (global)</code> <code>processname_VERBOSE = -2</code>	0

Lotus Notes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the client and click **Edit client**. Then click **Windows clients > Lotus Notes** or **UNIX client > Lotus Notes**.

The **Lotus Notes** properties apply to the currently selected client that runs NetBackup for Domino.

For more information, see the [NetBackup for HCL Domino Administrator's Guide](#).

For UNIX servers: If you have multiple installations of Domino server, the values in the client properties only apply to one installation. For other installations, specify the installation path and location of the `notes.ini` file with the `LOTUS_INSTALL_PATH` and `NOTES_INI_PATH` directives in the backup policy.

Table 2-43 Lotus Notes client host properties

Client host properties	Description
Maximum number of logs to restore	The maximum number of logs that can be prefetched in a single restore job during recovery. Specify a value greater than 1. A value less than or equal to 1, does not gather transaction logs during recovery. One transaction log extent per job is restored to the Domino server's log directory.
Transaction log cache path	A path where NetBackup can temporarily store the prefetched transaction logs during recovery. If you do not specify a path, during recovery NetBackup restores the logs to the Domino server's transaction log directory. Note the following: ■ If the specified path does not exist then it is created during restore. ■ The user must have write permission for the folder. ■ Transaction logs are restored to the original location, the Domino transaction log directory, if a path is not specified. ■ If the value of Maximum number of logs to restore is less than or equal to 1 then this path is ignored. The logs are not prefetched; one transaction log per job is restored to the Domino Server's log directory. ■ If there is not sufficient space to restore the specified number of logs, NetBackup tries to restore only the number of logs that can be accommodated.
INI path	The <code>notes.ini</code> file that is associated with the Domino partitioned servers used to back up and restore the Notes database. This setting does not apply to non-partitioned servers. ■ On Windows: If the <code>notes.ini</code> file is not located in the default directory, indicate its location. ■ On UNIX: If the <code>notes.ini</code> is not located in the directory that is specified in the Path, indicate its location here. Include the directory and the <code>notes.ini</code> file name.
Path	The path where the Notes program files reside on the client. NetBackup must know where these files are to perform backup and restore operations. ■ On Windows: The path for program directory (where <code>nserver.exe</code> resides). ■ On UNIX: A path that includes the Domino data directory, the Notes program directory, and the Notes resource directory.

Media properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server. If necessary, click **Connect**. Then click **Edit primary server** or **Edit media server**. Click **Media**.

The **Media** host properties contain the following settings.

Table 2-44 Media properties

Property	Description
Allow media overwrite property	This property overrides the NetBackup overwrite protection for specific media types. Normally, NetBackup does not overwrite certain media types. To disable overwrite protection, place a check in the check box of one or more of the listed media formats. For example, place a check in the CPIO check box to permit NetBackup to overwrite the cpio format. By default, NetBackup does not overwrite any of the formats on removable media, and logs an error if an overwrite attempt occurs. This format recognition requires that the first variable length block on a media be less than or equal to 32 kilobytes. The following media formats on removable media can be selected to be overwritten: ■ When ANSI is enabled, ANSI labeled media can be overwritten. ■ When TAR is enabled, TAR media can be overwritten. ■ When DBR is enabled, DBR media can be overwritten. (The DBR backup format is no longer used.) ■ Remote Storage MTF1 media format. When MTF1 is enabled, Remote Storage MTF1 media format can be overwritten. ■ When CPIO is enabled, CPIO media can be overwritten. ■ When AOS/VS is enabled, AOS/VS media can be overwritten. (Data General AOS/VS backup format.) ■ When MTF is enabled, MTF media can be overwritten. With only MTF checked, all other MTF formats can be overwritten. (The exception is Backup Exec MTF (BE-MTF1) and Remote Storage MTF (RS-MTF1) media formats, which are not overwritten. ■ When BE-MTF1 is enabled, Backup Exec MTF media can be overwritten. See “Results when media overwrites are not permitted” on page 122.

Table 2-44 Media properties (*continued*)

Property	Description
Enable SCSI reserve	This property allows exclusive access protection for tape drives. With access protection, other host bus adaptors cannot issue commands to control the drives during the reservation. SCSI reservations provide protection for NetBackup Shared Storage Option environments or any other multiple-initiator environment in which drives are shared. The protection setting configures access protection for all tape drives from the media server on which the option is configured. You can override the media server setting for any drive path from that media server. See “Recommended use for Enable SCSI reserve property” on page 123. See “Drive path options” on page 447. The following are the protection options: ■ The SCSI persistent reserve option provides SCSI persistent reserve protection for SCSI devices. The devices must conform to the SCSI Primary Commands - 3 (SPC-3) standard. ■ The SPC-2 SCSI reserve option (default) provides SPC-2 SCSI reserve protection for SCSI devices. The devices must conform to the reserve and release management method in the SCSI Primary Commands - 2 standard. ■ To operate NetBackup without tape drive access protection, clear the Enable SCSI reserve property. If unchecked, other HBAs can send the commands that may cause a loss of data to tape drives. Note: Ensure that all of your hardware processes SCSI persistent reserve commands correctly. All of your hardware includes Fibre Channel bridges. If the hardware does not process SCSI persistent reserve commands correctly and NetBackup is configured to use SCSI persistent reserve, no protection may exist.
Allow multiple retentions per media	This property lets NetBackup mix retention levels on tape volumes. It applies to media in both robotic drives and nonrobotic drives. The default is that the check box is clear and each volume can contain backups of only a single retention level.
Allow backups to span tape media	This property, when checked, lets backups span to multiple tape media. This property lets NetBackup select another volume to begin the next fragment. The resulting backup has data fragments on more than one volume. The default is that Allow backups to span tape media is checked and backups are allowed to span media. If the end of media is encountered and this property is not selected, the media is set to FULL and the operation terminates abnormally. This action applies to both robotic drives and nonrobotic drives.

Table 2-44 Media properties (*continued*)

Property	Description
Allow backups to span disk volumes	This property lets backups span disk volumes when one disk volume becomes full. The default is that this property is enabled. The Allow backups to span disk volumes property does not apply to AdvancedDisk or OpenStorage storage units. Backups span disk volumes within disk pools automatically. The following destinations support disk spanning: ■ A BasicDisk storage unit spanning to a BasicDisk storage unit. The units must be within a storage unit group. ■ An OpenStorage or AdvancedDisk volume spanning to another volume in the disk pool. For disk spanning to occur, the following conditions must be met: ■ The storage units must share the same media server. ■ The multiplexing level on spanning storage units should be the same. If there are any differences, the level on the target unit can be higher. See “Enable multiplexing storage unit setting” on page 561. ■ A disk staging storage unit cannot span to another storage unit. Also, a disk staging storage unit is not eligible as a target for disk spanning. ■ Disk spanning is not supported on NFS.
Enable standalone drive extension	This property lets NetBackup use whatever labeled or unlabeled media is found in a nonrobotic drive. The default is that the Enable standalone drive extension property is enabled.
Enable job logging	This property allows the logging of the job information. This logging is the same information that the NetBackup Activity monitor uses. The default is that job logging occurs.
Enable unrestricted media sharing for all media servers	This property controls media sharing, as follows: ■ Enable this property to allow all NetBackup media servers and NDMP hosts in the NetBackup environment to share media for writing. Do not configure server groups for media sharing. ■ Clear this property to restrict media sharing to specific server groups. Then configure media server groups and backup policies to use media sharing. ■ Clear this property to disable media sharing. Do not configure media server groups. The default is that media sharing is disabled. (The property is cleared and no server groups are configured.) See “About NetBackup server groups” on page 374.

Table 2-44 Media properties (*continued*)

Property	Description
Media ID prefix (non-robotic)	This property specifies the media ID prefix to use in media IDs when the unlabeled media is in nonrobotic drives. The prefix must be one to three alpha-numeric characters. NetBackup appends numeric characters. By default, NetBackup uses A and assigns media IDs such as A00000, A00001, and so on. For example, if FEB is specified, NetBackup appends the remaining numeric characters. The assigned media IDs become FEB000, FEB001, and so on.
Media unmount delay	To specify a Media unmount delay property indicates that the unloading of media is delayed after the requested operation is complete. Media unmount delay applies only to user operations, to include backups and restores of database agent clients, such as those running NetBackup for Oracle. The delay reduces unnecessary media unmounts and the positioning of media in cases where the media is requested again a short time later. The delay can range from 0 seconds to 1800 seconds. The default is 180 seconds. If you specify 0, the media unmount occurs immediately upon completion of the requested operation. Values greater than 1800 are set to 1800.
Media request delay (non-robotic)	This property specifies how long NetBackup waits for media in nonrobotic drives. During the delay period, NetBackup checks every 60 seconds to see if the drive is ready. If the drive is ready, NetBackup uses it. Otherwise, NetBackup waits another 60 seconds and checks again. If the total delay is not a multiple of 60, the last wait is the remainder. If the delay is less than 60 seconds, NetBackup checks after the end of the delay. For example, set the delay to 150 seconds. NetBackup waits 60 seconds, checks for ready, waits 60 seconds, checks for ready, waits 30 seconds, and checks for ready the last time. If the delay was 50 seconds (a short delay is not recommended), NetBackup checks after 50 seconds.

Results when media overwrites are not permitted

If media contains one of the protected formats and media overwrites are not permitted, NetBackup takes the following actions:

If the volume has not been previously assigned for a backup	■ Sets the volume's state to FROZEN ■ Selects a different volume ■ Logs an error
If the volume is in the NetBackup media catalog and was previously selected for backups	■ Sets the volume's state to SUSPENDED ■ Aborts the requested backup ■ Logs an error

If the volume is mounted for a backup of the NetBackup catalog	The backup is aborted and an error is logged. The error indicates the volume cannot be overwritten.
If the volume is mounted to restore files or list the media contents	NetBackup aborts the request and logs an error. The error indicates that the volume does not have a NetBackup format.

Recommended use for Enable SCSI reserve property

All tape drive and bridge vendors support the SPC-2 SCSI reserve and release method. NetBackup has used SPC-2 SCSI reserve since NetBackup 3.4.3, and it is the default tape drive reservation method in NetBackup. SPC-2 SCSI reserve is effective for most NetBackup environments.

The SCSI persistent reserve method provides device status and correction and may be more effective in the following environments:

- Where NetBackup media servers operate in a cluster environment.
NetBackup can recover and use a reserved drive after a failover (if NetBackup owns the reservation). (With SPC-2 SCSI reserve, the drive must usually be reset because the reservation owner is inoperative.)
- Where the drive has high availability.
NetBackup can resolve NetBackup drive reservation conflicts and maintain high drive availability. (SPC-2 SCSI reserve provides no method for drive status detection.)

However, the SCSI persistent reserve method is not supported or not supported correctly by all device vendors. Therefore, thoroughly analyze the environment to ensure that all of the hardware supports SCSI persistent reserve correctly.

It is recommended to carefully consider all of the following factors before **Enable SCSI reserve** is used:

- Only a limited number of tape drive vendors support SCSI persistent reserve.
- SCSI persistent reserve is not supported or not supported correctly by all Fibre Channel bridge vendors. Incorrect support in a bridge means no access protection. Therefore, if the environment uses bridges, do not use SCSI persistent reserve.
- If parallel SCSI buses are used, carefully consider the use of SCSI persistent reserve. Usually, parallel drives are not shared, so SCSI persistent reserve protection is not required. Also, parallel drives are usually on a bridge, and bridges do not support SCSI persistent reserve correctly. Therefore, if the environment uses parallel SCSI buses, do not use SCSI persistent reserve.

- The operating system tape drivers may require extensive configuration to use SCSI persistent reserve. For example, if the tape drives do not support SPC-3 Compatible Reservation Handling (CRH), ensure that the operating system does not issue SPC-2 reserve and release commands.

If any of the hardware does not support SCSI persistent reserve, it is not recommended that SCSI persistent reserve is used.

Network properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the client. If necessary click **Connect**, then click **Edit client**. Click **Windows client > Network**.

Use the **Network** properties to configure the communications requirements between clients and the primary server. These properties apply to the currently selected Windows client .

The **Network** host properties contain the following settings.

Table 2-45 Network properties for Windows clients

Property	Description
NetBackup client service port (BPCD)	This property specifies the port that the NetBackup client uses to communicate with the NetBackup server. The default is 13782. Note: If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.
NetBackup request service port (BPRD)	This property specifies the port for the client to use when it sends requests to the NetBackup request service (<code>bprd</code> process) on the NetBackup server. The default is 13720. Note: If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.
Announce DHCP interval	This property specifies how many minutes the client waits before it announces that a different IP address is to be used. The announcement occurs only if the specified time period has elapsed and the address has changed since the last time the client announced it.

Network settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Network settings**.

The **Network settings** host properties apply to primary servers, media servers, and clients.

The **Network settings** page contains properties for **Reverse host name lookup** and **Use the IP address family**.

See [“Reverse host name lookup property”](#) on page 125.

See [“Use the IP address family property”](#) on page 126.

Reverse host name lookup property

The domain name system (DNS) reverse host name lookup is used to determine what host and domain name a given IP address indicates.

Some administrators cannot or do not want to configure the DNS server for reverse host name lookup. For these environments, NetBackup offers the **Reverse host name lookup** property to allow, restrict, or prohibit reverse host name lookup.

Administrators can configure the **Reverse host name lookup** property for each host.

Table 2-46 Reverse host name lookup property settings

Property	Description
Allowed	The Allowed property indicates that the host requires reverse host name lookup to work to determine that the connection comes from a recognizable server. By default, the host resolves the IP address of the connecting server to a host name by performing a reverse lookup. If the conversion of the IP address to host name fails, the connection fails. Otherwise, it compares the host name to the list of known server host names. If the comparison fails, the host rejects the server and the connection fails.
Restricted	The Restricted property indicates that the NetBackup host first attempts to perform reverse host name lookup. If the NetBackup host successfully resolves the IP address of the connecting server to a host name (reverse lookup is successful), it compares the host name to the list of known server host names. If the resolution of the IP address to a host name fails (reverse lookup fails), based on the Restricted setting, the host converts the host names of the known server list to IP addresses (using a forward lookup). The host compares the IP address of the connecting server to the list of known server IP addresses. If the comparison fails, the host rejects the connection from server and the connection fails.

Table 2-46 Reverse host name lookup property settings (*continued*)

Property	Description
Prohibited	The Prohibited property indicates that the NetBackup host does not try reverse host name lookup at all. The host resolves the host names of the known server list to IP addresses using forward lookups. The NetBackup host then compares the IP address of the connecting server to the list of known server IP addresses. If the comparison fails, the NetBackup host rejects the connection from the server and the connection fails.

Reverse Host Name Lookup changes outside of the Administration Console

In some cases, a primary server may not be able to view the host properties of a media server or client in the **NetBackup Administration Console**. The NetBackup customer's DNS reverse host name lookup configuration may be one possible reason why the **Host Properties** may not be visible.

In this case, since changing the NetBackup **Reverse Host Name Lookup** host property requires being able to view the **Host Properties**, you'll need to use another method to change it.

Configure the `REVERSE_NAME_LOOKUP` option by using the `nbgetconfig` and `nbsetconfig` commands. The `nbsetconfig` command configures the option on Windows and UNIX primary servers and clients.

See [“Methods to set the NetBackup configuration options”](#) on page 48.

The `REVERSE_NAME_LOOKUP` entry uses the following format:

```
REVERSE_NAME_LOOKUP = ALLOWED | RESTRICTED | PROHIBITED
```

For example:

```
REVERSE_NAME_LOOKUP = PROHIBITED
```

The values of `ALLOWED`, `RESTRICTED`, and `PROHIBITED` represent the same meaning as the values in the **Network Settings** host properties.

Use the IP address family property

On the hosts that use both IPv4 and IPv6 addresses, use the **Use the IP address family** property to indicate which address family to use:

- **IPv4 only** (Default)
- **IPv6 only**

- Both IPv4 and IPv6

While the **Use the IP address family** property controls how host names are resolved to IP addresses, the **Preferred network** properties control how NetBackup uses the addresses.

Port ranges properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Port ranges**.

Use the**Port ranges** properties to determine how hosts connect to one another. These properties apply to the selected primary server, media server, or client.

The **Port ranges** host properties contain the following settings.

Table 2-47 Port ranges host properties

Property	Description
Use random port assignments	Specifies how the selected computer chooses a port when it communicates with NetBackup on other computers. Enable this property to let NetBackup randomly select ports from those that are free in the allowed range. For example, if the range is from 1023 through 5000, it chooses randomly from the numbers in this range. If this property is not enabled, NetBackup chooses numbers sequentially, not randomly. NetBackup starts with the highest number that is available in the allowed range. For example, if the range is from 1023 through 5000, NetBackup chooses 5000. If 5000 is in use, port 4999 is chosen. This property is enabled by default.
Client port window	Select Use OS selected non-reserved port to let the operating system determine which non-reserved port to use. Or, select the range of non-reserved ports on the selected computer. NetBackup can use any available port within this range as the source port when communicating with NetBackup on another computer.

Table 2-47 Port ranges host properties (*continued*)

Property	Description
Server port window	This property specifies the range of non-reserved ports on which NetBackup processes on this computer accept connections from NetBackup when the connection is not to a well known port. This property primarily applies to <code>bpcd</code> call-back when <code>vnetd</code> is disabled in the connect options and the local host name is configured for non-reserved ports. This property also applies in the situation where a third-party protocol is used, such as NDMP. It specifies the range of non-reserved ports on which this server accepts NetBackup connections from other computers. The default range is 1024 through 5000. Instead of indicating a range of ports, you can enable Use OS selected non-reserved port to let the operating system determine which non-reserved port to use. This setting applies to the selected primary or media server.
Server reserved port window	This entry specifies the range of local reserved ports on which this computer accepts connections from NetBackup when the connection is not to a well known port. This property primarily applies to <code>bpcd</code> call-back when <code>vnetd</code> is disabled in the connect options for a local host name. Instead of indicating a range of ports, you can enable Use OS selected non-reserved port to let the operating system determine which non-reserved port to use.

Registered ports and dynamically-allocated ports

NetBackup communicates between computers by using a combination of registered ports and dynamically-allocated ports.

Registered ports

These ports are registered with the Internet Assigned Numbers Authority (IANA) and are permanently assigned to specific NetBackup services. For example, the port for the NetBackup client daemon (`bpcd`) is 13782.

The following system configuration file can be used to override the default port numbers for each service:

On Windows: `%systemroot%\system32\drivers\etc\services`

On UNIX: `/etc/services`

Note: It is not recommended to change the port numbers that are associated with PBX (1556 and 1557).

Dynamically-allocated ports

These ports are assigned as needed, from configurable ranges in the **Port ranges** host properties for NetBackup servers and clients.

In addition to the range of numbers, you can specify whether NetBackup selects a port number at random or starts at the top of the range and uses the first one available.

Preferred network properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Preferred network**.

Use the **Preferred network** properties to specify to NetBackup which networks or interfaces to use for outgoing NetBackup traffic from the selected hosts. These properties apply to currently selected primary server, media server, or client.

Note: The Preferred network setting in NetBackup does not apply to the Granular Recovery Technology (GRT) and VMware Instant Recovery features. Network settings that are configured in the operating system are used for these features during communication.

Preferred network entries are not needed if NetBackup is configured using host names with IP addresses to which the operating system resolves and then routes correctly.

When external constraints prevent the environment from being corrected, **Preferred network** entries may be useful as follows:

- Can be used to prevent NetBackup from connecting to specific destination addresses.
- Can be used to cause NetBackup to connect only to specific destination addresses.
- Can be used to request a subset of local interfaces for source binding when making outbound connections.

Caution: When used for source binding, the operating system may not honor the source binding list provided by NetBackup. If the operating system implements the weak host model, asymmetrical network routing may result. If asymmetrical routing occurs, the remote host may reject the inbound connection if it implements the strong host model. Similarly, stateful network devices may also drop asymmetrical connections. To ensure the use of specific outbound interfaces for specific remote hosts or networks, make sure that the OS name resolution and routing configurations are correct; create static host routes if needed. Ensure that all network drivers properly implement the IP and TCP networking protocols.

The local **Preferred network** entries do not affect the forwarding profile that the local host returns to a remote host during initial CORBA connection setup; it contains all the local plumbed interfaces. However, the End Point Selection algorithm within the remote process uses its local **Preferred network** entries to evaluate the profile when it selects the destination for the subsequent CORBA connection.

With respect to source binding, the **Preferred network** properties offer more flexibility than the **Use specified network interface** property in the **Universal settings** properties. The **Use specified network interface** property can be used to specify only a single interface for NetBackup to use for outbound calls. The **Preferred network** properties were introduced so that administrators can give more elaborate and constrictive instructions that apply to multiple individual networks, or a range of networks. For example, an administrator can configure a host to use any network except one. If both properties are specified, **Use specified network interface** overrides **Preferred network**.

Note: Do not inadvertently configure hosts so that they cannot communicate with any other host. Use the `bptestnetconn` utility to determine whether the hosts can communicate as you intend.

See [“bptestnetconn utility to display Preferred network information”](#) on page 140.

The **Preferred network** host properties contain a list of networks and the directive that has been configured for each.

Table 2-48 Preferred network host properties

Property	Description
List of network specifications for NetBackup communications	The list of preferred networks contains the following information: ■ The Target column lists the networks (or host names or IP addresses) that have been given specific directives. If a network is not specifically listed as a target, or if a range of addresses does not include the target, NetBackup considers the target to be available for selection. Note that if the same network considerations apply for all of the hosts, the list of directives can be identical across all hosts in the NetBackup environment. If a directive contains an address that does not apply to a particular host, that host ignores it. For example, an IPv4-only host ignores IPv6 directives, and IPv6-only hosts ignore IPv4 directives. This action lets the administrator use the same Preferred network configurations for all the hosts in the NetBackup environment. ■ The Specified as column indicates the directive for the network: Match, Prohibited, or Only. ■ The Source column lists source binding information to use to filter addresses. The Source property is an optional configuration property.
Ordering arrows	Select a network in the list, then click the up or down arrow to change the order of the network in the list. The order can affect which network NetBackup selects. See “Order of directive processing in the Preferred network properties” on page 139.
Add	Click Add to add a network to the Preferred network properties. Then configure the directive for the network. See Table 2-49 on page 132.
Actions > Edit	Locate a network in the list, then click Actions > Edit to change the Preferred network properties. See “Add or edit a Preferred network setting” on page 131.
Actions > Delete	Locate a network in the list, then click Actions > Delete to remove the network from the list of preferred networks.

Add or edit a Preferred network setting

Refer to the following settings when you add or edit a preferred network setting.

Table 2-49 Configuration for Preferred network settings

Property	Description
Target	Enter a network address or a host name: - NetBackup recognizes the following wildcard entries as addresses: <code>0.0.0.0</code> Matches any IPv4 address. <code>0::0</code> Matches any IPv6 address. <code>0/0</code> Matches the address of any family. - If the target is a host name which resolves to more than one IP address, only the first IP address will be used. - If a subnet is not specified, the default is /128 when the address is non-zero and /0 when the address is 0. This applies to both Target and Source properties. A subnet of /0 cannot be used with a non-zero address because it effectively negates all of the bits in the address, making the target or the source match every address. For example, 0/0. Note: Do not use the following malformed entries as wildcards: 0/32, 0/64, or 0/128. The left side of the slash must be a legitimate IP address. However, 0/0 may be used, as listed.
Match	The Match directive: - Applies when Target is a destination address. - Indicates that the specified network, address, or host name is preferred for communication with the selected host. - Does not reject other networks, addresses, or host names from being selected, even if they do not match. (The Only directive rejects unsuitable targets if they do not match.) - Is useful following a Prohibited or a Only directive. When used with other directives, Match indicates to NetBackup to stop rule processing because a suitable match has been found. - Can be used with the Source property to indicate source binding.
Prohibited	Use the Prohibited directive to exclude or prevent the specified network, address, or host name from being used. The Target is applied to both the source and the destination addresses. If a Source is specified and the Prohibited is indicated, it is ignored but the target is still prohibited. If the matched address is a destination address, evaluation stops. If this was the only potential destination, the connection is not attempted. If there are additional potential destinations, they are evaluated starting over with the first entry. If the matched address is a source address, it is removed from the source binding list. Caution: On some platforms, prohibiting a local interface may cause unexpected results when connecting to remote hosts. Prohibiting a local interface does not affect connections that are internal to the host.

Table 2-49 Configuration for Preferred network settings (*continued*)

Property	Description
Only	The Only directive: ■ Applies to destination addresses. ■ Indicates that the specified network, address, or host name that is used for communication with the selected host must be in the specified network. Use the Only directive to prevent any network from being considered other than those specified as Only. ■ If the address that is being evaluated does not match the target, it is not used and evaluation stops for that address. If the address being evaluated was the only potential destination, the connection is not attempted. If there is an additional potential destination, it is evaluated starting over with the first entry. ■ Can be used with the Source property to indicate source binding.
Source	Use this property with the Match or the Only directives to identify the local host name, IP addresses, or networks that may be used for source binding. If a subnet is not specified, the default is /128. If this host has an IP address that matches Source, that IP address will be used as the source when connecting to the destination. If the Source is not valid for this host, it is ignored.

How NetBackup uses the directives to determine which network to use

Each host has an internal table of preferred network rules that NetBackup consults before it selects a network interface to use for communication with another host. The table includes every interface-IP address combination available to the selected host. Based on the **Preferred NetBackup** directives, the table indicates to NetBackup whether or not the host is allowed to use a given network.

This topic uses the example of two multihomed servers (Server_A and Server_B) as shown in [Figure 2-3](#). Server A is considering which addresses it can use to access Server_B, given the **Preferred network** directives configured on Server_A.

When **Preferred network** directives are used to place restrictions on targets, they are added from the perspective of the server making the connection. The directives on Server_A affect its preferences as to which Server_B addresses it can use.

Figure 2-3 Multihomed servers example




Figure 2-4 shows a table for Server_B. Server_B has multiple network interfaces, some of which have multiple IP addresses. In the table, *yes* indicates that NetBackup can use the network-IP combination as a source. In this example, no directives have been created for the host. Since no networks are listed in the **Preferred network** properties, any network-IP combinations can be used for communication.

Note: The following topic shows the `bptestnetconn` output for this example configuration:

See [“bptestnetconn utility to display Preferred network information”](#) on page 140.

Figure 2-4 From Server_A's perspective: Available IP addresses on Server_B when no directives are indicated on Server_A

IP addresses		
Network interfaces		
	IPv4	IPv6
	2001:0db8:0:1f0::1efc	---
	10.80.73.147	Yes
	2001:0db8:0:11c::1efc	---
	2001:0db8:0:11d::1efc	Yes
	2001:0db8:0:11e::1efc	Yes
	10.96.73.253	Yes

Figure 2-5 shows a table for the same host (Server_B). Now, the **Preferred network** properties are configured so that all IPv4 addresses are excluded from selection consideration by NetBackup. All NetBackup traffic is to use only IPv6 addresses.

Figure 2-5 From Server_A's perspective: Available IP addresses on Server_B when directives to use IPv6 addresses only are indicated on Server_A

IP addresses		
Network interfaces		
	IPv4	IPv6
	2001:0db8:0:1f0::1efc	---
	10.80.73.147	No
	2001:0db8:0:11c::1efc	---
	2001:0db8:0:11d::1efc	Yes
	2001:0db8:0:11e::1efc	Yes
	10.96.73.253	No

The following topics describe various configurations:

- See [“Configurations to use IPv6 networks”](#) on page 136.
- See [“Configurations to use IPv4 networks”](#) on page 138.
- See [“Configuration to prohibit using a specified address”](#) on page 141.
- See [“Configuration to prefer a specified address”](#) on page 142.
- See [“Configuration that restricts NetBackup to one set of addresses”](#) on page 143.

- See [“Configuration that limits the addresses, but allows any interfaces”](#) on page 144.

Configurations to use IPv6 networks

The following **Preferred network** configurations instruct NetBackup to use only IPv6 addresses as targets in outbound calls for the currently selected hosts. The configurations satisfy a topology where all backup traffic uses an IPv6 network and other traffic uses other networks.

One configuration uses the **Prohibited** directive ([Figure 2-6](#)) and one configuration uses the **Match** directive ([Figure 2-7](#)).

The more efficient method to specify one address family, (IPv6, in this case), is to prohibit IPv4. The behavior of the **Match** directive is not as exclusive as **Prohibited**. In this case, **Match** may not necessarily exclude other address families.

[Figure 2-6](#) uses the **Prohibited** directive with a wildcard to indicate to NetBackup to not consider using any IPv4 addresses. In this situation, NetBackup must use an IPv6 address.

Note: The default configuration is for NetBackup to use only IPv4 addresses.

If you have not previously changed the **Network settings > Use the IP address family** option to **Both IPv4 and IPv6** or **IPv6 only**, creating a directive that prohibits all IPv4 addresses renders the server mute.

See [“Use the IP address family property”](#) on page 126.

See [“Network settings properties”](#) on page 124.

Figure 2-6 Prohibit IPv4 addresses as targets

Add preferred network settings

Target
0.0.0.0

Specified as

☐ Match (The above network is preferred for communication)

☒ Prohibited (The above network is not used for communication)


☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Figure 2-7 uses the **Match** directive with a wildcard to indicate to NetBackup to prefer IPv6 addresses. In this case, NetBackup tries to use an IPv6 address, but may consider IPv4 addresses if necessary.

Figure 2-7 Match IPv6 addresses as targets



The screenshot shows a dialog box titled "Add preferred network settings" with a close button (X) in the top right corner. Inside the dialog, there is a "Target" section with the text "0::0" below it. Below the "Target" section is a "Specified as" section with three radio button options: "Match (The above network is preferred for communication)" (which is selected), "Prohibited (The above network is not used for communication)", and "Only (Only target addresses in the above network is used for communication)". Below the "Specified as" section is a "Source" section with an empty text field. At the bottom of the dialog are three buttons: "Cancel", "Add and add another", and "Add".

Figure 2-8 shows another configuration that allows NetBackup to choose from multiple IPv6 networks.

Given the multihomed example configuration, the directive indicates the following:

- Four IPv6 networks, from `fec0:0:0:fe04` through `fec0:0:0:fe07`, are described as targets.
- For all addresses in these networks, a source binding address that is derived from the IP addresses of host name `host_fred` is used.

See [“How NetBackup uses the directives to determine which network to use”](#) on page 133.

Figure 2-8 Indicating a range of IPv6 networks

Add preferred network settings

Target
fec0:0:0:fe04::/62

Specified as

☐ Match (The above network is preferred for communication)

☐ Prohibited (The above network is not used for communication)

☒ Only (Only target addresses in the above network is used for communication)

Source
host_fred

Cancel Add and add another Add


Configurations to use IPv4 networks

The following **Preferred network** configurations instruct NetBackup to use only IPv4 addresses as targets in outbound calls for the currently selected hosts. The configurations satisfy a topology where all backup traffic uses an IPv4 network and other traffic uses other networks.

One configuration uses the **Prohibited** directive (Figure 2-9) and one configuration uses the **Match** directive (Figure 2-10).

The more efficient method to specify one address family, (IPv4, in this case), is to prohibit IPv6. The behavior of the **Match** directive is not as exclusive as **Prohibited**. In this case, **Match** may not necessarily exclude other address families.

Figure 2-9 uses the **Prohibited** directive with a wildcard to indicate to NetBackup to not consider using any IPv6 addresses. In this situation, NetBackup must use an IPv4 address.

Figure 2-9 Prohibit IPv6 addresses as targets

Add preferred network settings

Target
0::0

Specified as

☐ Match (The above network is preferred for communication)


☒ Prohibited (The above network is not used for communication)

☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

[Figure 2-10](#) uses the **Match** directive with a wildcard to indicate to NetBackup to prefer IPv4 addresses. In this case, NetBackup tries to use an IPv4 address, but may consider IPv6 addresses if necessary.

Figure 2-10 Match IPv4 addresses as targets

Add preferred network settings

Target
0.0.0.0

Specified as

☒ Match (The above network is preferred for communication)

☐ Prohibited (The above network is not used for communication)

☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Order of directive processing in the Preferred network properties

NetBackup sorts all directives into decreasing order by the **Target** subnet length so that the more specific network specifications, such as complete host names or IP addresses, match first. (For example, a **Target** with a /24 subnet is processed before a **Target** with a /16 subnet.) In this way, NetBackup can honor host-specific overrides.

If multiple directives have the same length subnet, NetBackup looks at the order in which the directives are listed.

Use the up or down arrows to the right of the list to change the order of the directives.

NetBackup processes each resolved destination address and each prospective source address relative to the directives. Directives that contain addresses that do not apply to either host are ignored.

bptestnetconn utility to display Preferred network information

The `bptestnetconn` utility is available to administrators to test and analyze host connections. Use the preferred network option (`--prefnet` or `-p`) to display information about the preferred network configuration, along with the forward lookup information of a host on the server list.

For example, `bptestnetconn -v6 -p -s -H host1` displays the directives in the order in which NetBackup processes them, which may not be the order in which they are configured.

- The `bptestnetconn` command is described in the [NetBackup Commands Reference Guide](#).
- The following article contains best practices for using `bptestnetconn` command: https://www.veritas.com/content/support/en_US/article.100009286

Figure 2-11 shows the `bptestnetconn` output when run on Server_A, for Server_B. That is, `bptestnetconn` is run from Server_A's perspective. Based on the directives configured on Server_A, for Server_B, `bptestnetconn` shows the available IP addresses on Server_B. In this example, no directives are configured on Server_A.

Figure 2-11 `bptestnetconn` for Server_B with no directives listed

```
[root@Server_A netbackup]# bptestnetconn -f --prefnet -H Server_B
-----
FL: Server_B -> 10.81.73.147           : 11 ms SRC: ANY
FL: Server_B -> 10.96.73.253          : 11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11d::1efc   : 11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11e::1efc   : 11 ms SRC: ANY
FL: Server_B -> 2001:d8b:0:1f0::1efc   : 11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11c::1efc   : 11 ms SRC: ANY
-----
Total elapsed time: 0 sec
```

Host for which lookup is performed	List of networks available to Server_B	Any source is available to use for a connection

The following directive is added to the **Preferred network** properties on Server_A:

In the configuration file the directive appears as follows:

```
PREFERRED_NETWORK = 2001:0db8:0:11c::/62 ONLY
```

This directive provides NetBackup with the information to filter the addresses and choose to communicate with only those that match the :11c, :11d, :11e, and :11f networks. The addresses that do not match the **Only** directive are prohibited, as shown in the `bptestnetconn` output.

Figure 2-12 shows the `bptestnetconn` output for Server_B, given this directive.


Figure 2-12 `bptestnetconn` for Server_B with directive

```
[root@Server_A netbackup]# bptestnetconn -f --prefnet -H Server_B
-----
FL: Server_B -> 10.81.73.147          :      11 ms TGT PROHIBITED
FL: Server_B -> 10.96.73.253          :      11 ms TGT PROHIBITED
FL: Server_B -> 2001:db8:0:11d::1efc   :      11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11e::1efc   :      11 ms SRC: ANY
FL: Server_B -> 2001:d8b:0:1f0::1efc   :      11 ms TGT PROHIBITED
FL: Server_B -> 2001:db8:0:11c::1efc   :      11 ms SRC: ANY
-----
Total elapsed time: 0 sec
```

List of networks available to Server_B
Directives make some targets unavailable to Server_B

Configuration to prohibit using a specified address

Figure 2-13 shows a configuration that prohibits NetBackup from using the specified address, or in this case, addresses.

Figure 2-13 Prohibited target example

Add preferred network settings

Target
192.168.100.0/24

Specified as

☐ Match (The above network is preferred for communication)

☒ Prohibited (The above network is not used for communication)

☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Configuration to prefer a specified address

Figure 2-14 shows a configuration that makes NetBackup prefer to use one range of destination addresses over others that might be available.

Other available destination addresses will only be used if one of the following is true:

- No destination address exists that is in this range, or
- A **Match** is specified for those addresses using a larger subnet mask, or
- A **Match** is specified for those addresses with a same length subnet mask and the address is ordered above this directive.

A **Prohibited** directive can be used to prevent the use of an address within this range. The **Prohibited** directive would need either a longer subnet mask, or a subnet mask of equal length with the **Prohibited** directive ordered above the **Match** directive. Additional **Match** directives may be used to indicate the additional backup networks that are allowed.

Figure 2-14 Match network selection with the source

Add preferred network settings

Target
192.168.100.0/24

Specified as

- ☒ Match (The above network is preferred for communication)
- ☐ Prohibited (The above network is not used for communication)
- ☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Configuration that restricts NetBackup to one set of addresses

Figure 2-15 configures NetBackup to use only the specified range of destination addresses, and the allowed source addresses must also be in the same range. The only exception is if other directives with larger subnets are present, or with equal-length subnets but ordered above this one.

Figure 2-15 Only network selection with the same source binding address

Add preferred network settings

Target
192.168.100.0/24

Specified as

- ☐ Match (The above network is preferred for communication)
- ☐ Prohibited (The above network is not used for communication)
- ☒ Only (Only target addresses in the above network is used for communication)

Source
192.168.100.0/24

Cancel Add and add another Add

A host with the **Only** directive configured considers only those target addresses in the 192.168.100.0 subnet. Additionally, source binding to the local interface must be done on the 192.168.100.0 subnet.

Configuration that limits the addresses, but allows any interfaces

Figure 2-16 shows a configuration that allows only the addresses that start with the specified prefix to be considered. No source binding is specified, so any interface may be used.

Figure 2-16 Limiting the addresses, without any source binding

The screenshot shows a dialog box titled "Add preferred network settings". It has a close button (X) in the top right corner. The "Target" field contains the text "fec0:0:1::/48". Below this, under the heading "Specified as", there are three radio buttons: "Match (The above network is preferred for communication)", "Prohibited (The above network is not used for communication)", and "Only (Only target addresses in the above network is used for communication)". The "Only" radio button is selected. Below the radio buttons is a "Source" field, which is currently empty. At the bottom of the dialog, there are three buttons: "Cancel", "Add and add another", and "Add".

Properties setting in host properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Properties**.

The host property **Properties** includes the following information about the selected host.

Table 2-50 Properties information for a host

Property name	Description
Host	The NetBackup client name of the host.
Operating system	The operating system and OS version on which the host is installed.
OS type	The type of OS.
Host type	The type of host: Primary server, media server, or client.
IP address	The IP address of the host.

RHV access hosts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **RHV access hosts**.

You can also configure these settings in the web UI from **Workloads > RHV**. Then select **RHV settings > Access hosts**.

Use the **RHV access hosts** to configure a special host that is called an RHV access host. This host is a NetBackup client that performs backups on behalf of the virtual machines. These properties apply to the currently selected primary server.

For more information, see the [NetBackup Red Hat Virtualization Administrator's Guide](#).

Resilient network properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Resilient network**.

For media servers and clients, the **Resilient network** properties are read only. When a job runs, the primary server updates the media server and the client with the current properties.

The **Resilient network** properties let you configure NetBackup to use resilient network connections for backups and restores. A resilient connection allows backup and restore traffic between a client and a NetBackup media server to function effectively in high-latency, low-bandwidth networks such as WANs. The data travels across a wide area network (WAN) to media servers in a central datacenter.

NetBackup monitors the socket connections between the remote client and the NetBackup media server. If possible, NetBackup re-establishes dropped connections and resynchronizes the data stream. NetBackup also overcomes latency issues to maintain an unbroken data stream. A resilient connection can survive network interruptions of up to 80 seconds. A resilient connection may survive interruptions longer than 80 seconds.

The NetBackup Remote Network Transport Service manages the connection between the computers. The Remote Network Transport Service runs on the primary server, the client, and the media server that processes the backup or restore job. If the connection is interrupted or fails, the services attempt to re-establish a connection and synchronize the data.

NetBackup protects only the network socket connections that the NetBackup Remote Network Transport Service (`nbrntd`) creates. Examples of the connections that are not supported are:

- Clients that back up their own data (deduplication clients and SAN clients)
- Granular Recovery Technology (GRT) for Exchange Server or SharePoint Server
- NetBackup `nbfsd` process.

NetBackup protects connections only after they are established. If NetBackup cannot create a connection because of network problems, there is nothing to protect.

Resilient connections apply between clients and NetBackup media servers, which includes primary servers when they function as media servers. Resilient connections do not apply to primary servers or media servers if they function as clients and back up data to a media server.

Resilient connections can apply to all of the clients or to a subset of clients.

Note: If a client is in a subdomain that is different from the server subdomain, add the fully qualified domain name of the server to the client's hosts file. For example, `india.domain.org` is a different subdomain than `china.domain.org`.

When a backup or restore job for a client starts, NetBackup searches the **Resilient network** list from top to bottom looking for the client. If NetBackup finds the client, NetBackup updates the resilient network setting of the client and the media server that runs the job. NetBackup then uses a resilient connection.

Table 2-51 Resilient network properties

Property	Description
FQDN or IP address	The full qualified domain name or IP address of the host. The address can also be a range of IP addresses so you can configure more than one client at once. You can mix IPv4 addresses and ranges with IPv6 addresses and subnets. If you specify the host by name, it is recommended that you use the fully qualified domain name. Use the arrow buttons on the right side of the pane to move up or move down an item in the list of resilient networks.
Resiliency	Resiliency is either On or Off .

Note: The order is significant for the items in the list of resilient networks. If a client is in the list more than once, the first match determines its resilient connection status. For example, suppose you add a client and specify the client IP address and specify **On** for **Resiliency**. Suppose also that you add a range of IP addresses as **Off**, and the client IP address is within that range. If the client IP address appears before the address range, the client connection is resilient. Conversely, if the IP range appears first, the client connection is not resilient.

Other NetBackup properties control the order in which NetBackup uses network addresses.

The NetBackup resilient connections use the SOCKS protocol version 5.

Resilient connection traffic is not encrypted. It is recommended that you encrypt your backups. For deduplication backups, use the deduplication-based encryption. For other backups, use policy-based encryption.

Resilient connections apply to backup connections. Therefore, no additional network ports or firewall ports must be opened.

Note: If multiple backup streams run concurrently, the Remote Network Transport Service writes a large amount of information to the log files. In such a scenario, it is recommended that you set the logging level for the Remote Network Transport Service to 2 or less. Instructions to configure unified logs are in a different guide.

View the resiliency status of a client

You can view the resiliency status of a client on the **Clients** tab of a policy or in the host properties for a client.

See [“Resilient network properties”](#) on page 145.

To view the resiliency status of a client in a policy

- 1 In the **NetBackup web UI**, open a policy.
- 2 Select the **Clients** tab.
- 3 The **Resiliency** column shows the status for each client in the policy.

To view the resiliency status of a client in host properties

- 1 In the **NetBackup web UI**, select **Host > Host properties**.
- 2 Select the client. If necessary, click **Connect**, then click **Edit client**.
- 3 Select **Resilient network**.

The **Resiliency** column shows the status for the client.

About Resilient jobs

The Resilient jobs feature lets the media server's job processes continue to run during a service disruption with the primary server. Backup metadata is cached to a user-defined location while the primary server processes are disrupted. Once the primary server re-establishes connections to the active media server processes, the cached data is transferred, and the backup proceeds.

To determine if a job is resilient, search the job details for the text, "job is resilient". If this text is present, the job is resilient.

The Resilient jobs feature is enabled by default. This feature is only available for some policy types. Please review the current requirements and limitations:

- The resiliency feature is either enabled or disabled. Backup jobs run as resilient jobs only when resiliency is enabled.
- Resilient jobs are only supported for Windows and Standard policy types.
- Backups cannot be multiplexed.
- Backups cannot have parent and child hierarchy. Use the Activity monitor to show parent and child relationship.
- Resilient jobs support the failure of the primary server. If the media server fails for any reason, the resilient jobs feature is not supported.

Note: If the primary server is also either the media server or the client, and it fails, the job is not resilient.

- If the client fails for any reason, the resilient job feature is not supported.
- If the primary server is upgraded while a backup is active, the backup is not resilient.
- The media server must be at NetBackup version 10.1.1 or later.
- Multistreamed backup jobs are not supported.
- Fiber Transport Media Server (FTMS) environments are not supported.

Resilient connection resource usage

Resilient connections consume more resources than regular connections, as follows:

- More socket connections are required per data stream. Three socket connections are required to accommodate the Remote Network Transport Service that runs on both the media server and the client. Only one socket connection is required for a non-resilient connection.

- More sockets are open on media servers and clients. Three open sockets are required rather than one for a non-resilient connection. The increased number of open sockets may cause issues on busy media servers.
- More processes run on media servers and clients. Usually, only one more process per host runs even if multiple connections exist.
- The processing that is required to maintain a resilient connection may reduce performance slightly.

Specify resilient connections for clients

Use the following procedure to specify resilient connections for NetBackup clients.

See [“Resilient network properties”](#) on page 145.

Alternatively, you can use the `resilient_clients` script to specify resilient connections for clients:

- Windows: `install_path\NetBackup\bin\admincmd\resilient_clients`
- UNIX: `/usr/opensv/netbackup/bin/admincmd/resilient_clients`

To specify resilient connections for clients

- 1 Open the **NetBackup web UI**.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Resilient network**.
- 5 You can perform the following actions:

Add a setting

To add a host or IP address setting

- 1 Click **Add**.
- 2 Enter a client host name or an IP address.

If you specify the client host by name, it is recommended that you use the fully qualified domain name.
- 3 Ensure that the **On** option is selected.
- 4 Click **Add and add another**.
- 5 Repeat until you have added each setting.
- 6 When you finish adding network settings, click **Add**.

- Edit a setting

To edit a host or IP address setting

1

Locate the client host name or the IP address.

2

Click **Actions > Edit**.

3

Select the desired **Resiliency** setting.

4

Click **Save**.
- Delete a setting

Delete a host or IP address setting

1

Locate the client host name or the IP address.

2

Click **Actions > Delete**.
- Up arrow,
Down arrow

Change the order of items

1

Select the client host name or the IP address.

2

Click the **Up** or **Down** button.

The order of the items in the list is significant.

See ["Resilient network properties"](#) on page 145.
- 6

If you want to begin a backup immediately, restart the NetBackup services on the primary server.

Resource limit properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Resource limits**.

The **Resource limits** properties control the number of simultaneous backups that can be performed on a particular resource type. These settings apply to all policies for the currently selected primary server.

Note: The **Resource limit** properties apply only to policies that use automatic selection of virtual machines (the policy's Query Builder). If you select virtual machines manually, the **Resource limit** properties have no effect.

See the respective guide for the workload or agent for details on the available resource limit properties.

Restore failover properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Restore failover**.

The **Restore failover** properties control how NetBackup performs automatic failover to a NetBackup media server. A failover server may be necessary if the regular media server is temporarily inaccessible to perform a restore operation. The automatic failover does not require administrator intervention. By default, NetBackup does not perform an automatic failover. These properties apply to currently selected primary servers.

The **Restore failover** host properties contain the following settings.

Table 2-52

Property	Description
Media server	Displays the NetBackup media servers that have failover protection for restores.
Failover restore servers	Displays the servers that provide the failover protection. NetBackup searches from top to bottom in the column until it finds another server that can perform the restore.

A NetBackup media server can appear only once in the **Media server** column but can be a failover server for multiple other media servers. The protected server and the failover server must both be in the same primary and media server cluster.

The following situations describe examples of when to use the restore failover capability:

- Two or more media servers share a robot and each has connected drives. When a restore is requested, one of the servers is temporarily inaccessible.
- Two or more media servers have standalone drives of the same type. When a restore is requested, one of the servers is temporarily inaccessible.

In these instances, inaccessible means that the connection between `bprd` on the primary server and `bptm` on the media server (through `bpcd`) fails.

Possible reasons for the failure are as follows:

- The media server is down.
- The media server is up but `bpcd` does not respond. (For example, if the connection is refused or access is denied.)

- The media server is up and `bpcd` is running, but `bptm` has problems. (For example, `bptm` cannot find the required tape.)

Assigning an alternate media server as a failover restore server

You can assign another media server to act as a failover restore server for your media server. If your media server is unavailable during a restore, the failover restore server takes its place.

To assign an alternate media server as a failover restore server

- 1 In the **NetBackup web UI** click **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Restore failover**.
- 5 Click **Add**.
- 6 In the **Media server** field, specify the media server for failover protection.
- 7 In the **Failover restore servers** field, specify the media servers to try if the server that is designated in the **Media server** field is unavailable. Separate the names of multiple servers with a single space.
- 8 Click **Add**.
- 9 Click **Save**.

Before the change takes effect, you must stop and restart the NetBackup Request Daemon on the primary server where the configuration was changed.

See [“About enabling automatic failover to an alternate server”](#) on page 1109.

Retention periods properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Retention periods**.

Use the **Retention periods** properties to define a duration for each retention level. You can select from 0-100 retention levels.

In a policy, the retention period determines how long NetBackup retains the backups or the archives that are created according to the schedule. These properties apply to selected primary servers.

By default, NetBackup stores each backup on a volume that already contains backups at the same retention level. However, NetBackup does not check the

retention period that is defined for that level. When the retention period for a level is redefined, some backups that share the same volume may have different retention periods.

For example, if the retention level 3 is changed from one month to 6 months, NetBackup stores future level 3 backups on the same volumes. That is, the backups are placed on the volumes with the level 3 backups that have a retention period of one month.

No problem exists if the new and the old retention periods are of similar values. However, before a major change is made to a retention period, suspend the volumes that were previously used for that retention level.

See [“Determining retention periods for volumes”](#) on page 155.

See [“Suspending or unsuspending volumes”](#) on page 520.

The **Retention periods** host properties contain the following settings.

Table 2-53 Retention periods page properties

Property	Description
Retention level	The retention level number (0 through 100). Value Assigns a number to the retention level setting. Units Specifies the units of time for the retention period. The list includes hours as the smallest unit of granularity and the special units, Infinite, and Expires immediately.
Retention period	A list of the current definitions for the possible levels of retention. By default, levels 9 through 100 (except level 25) are set to infinite. Retention level 9 cannot be changed and the retention period is always set to infinite. Retention level 25 also cannot be changed and the retention period is always set to expire immediately. With the default, there is no difference between a retention level of 12 and a retention level of 20, for example. If the retention period is changed for a level, it affects all schedules that use that level. The Changes pending column uses an asterisk (*) to indicate that the period has been changed and not applied. NetBackup does not change the actual configuration until the administrator accepts or applies the changes.
Schedule count	Lists the number of schedules that use the currently selected retention level.
Changes pending	This column displays an asterisk (*) to indicate that the period has been changed and not applied. NetBackup does not change the actual configuration until the administrator accepts or applies the changes.

Table 2-53 Retention periods page properties (*continued*)

Property	Description
Schedules using this retention level	Displays a list of the current policy names and schedule names that use the retention level.
Impact report	Displays a summary of how changes affect existing schedules. The list displays all schedules in which the retention period is shorter than the frequency period.

Changing a retention period

Use the following procedure to change a retention period.

To change a retention period

- 1 Open the web UI.
- 2 On the left, select **Hosts > Host properties**.
- 3 Select the primary server.
- 4 If necessary, click **Connect**. Then click **Actions > Edit primary server**.
- 5 Click **Retention periods**.
- 6 Locate the retention level to change and click **Edit**.

By default, levels 9 through 100 (except level 25) are set to infinite. If the levels are left at the default, there is no difference between a retention level of 12 and a retention level of 20. Level 9 cannot be changed and the retention period is always set to infinite. Retention level 25 also cannot be changed and the retention period is always set to expires immediately.

NetBackup displays the names of all schedules that use the selected retention level as well as the policy to which each schedule belongs.

- 7 Type the new retention period in the **Value** box.

- 8 From the **Units** drop-down list, select a unit of measure (days, weeks, months, years, infinite, or expires immediately).

After you change the value or unit of measure, an asterisk (*) appears in the **Changes pending** column to indicate that the period was changed. NetBackup does not change the actual configuration until the administrator accepts or applies the changes.

- 9 Click **Impact report**.

The policy impact list displays the policies and the schedule names where the new retention period is less than the frequency period. To prevent a potential gap in backup coverage, redefine the retention period for the schedules or change the retention or frequency for the schedule.

Determining retention periods for volumes

Use the following procedure to determine retention periods for volumes.

To determine retention periods for volumes

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**.
- 3 Click the **Volumes** tab. Find the volume in the list and examine the value in the **Retention period** column.

To see all volumes that have the same retention period, click the **Retention period** column header to sort the volumes by retention period.

Scalable Storage properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the media server. If necessary click **Connect**, then click **Edit media server**. Click **Scalable storage**.

The **Scalable Storage** properties contain information about encryption, metering, bandwidth throttling, and network connections between the NetBackup hosts and your cloud storage provider. These properties appear only if the host is supported for cloud storage. See the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List* for your release available through the following URL:

https://www.veritas.com/support/en_US/article.100040093

The **Scalable storage** properties apply to currently selected media server .

The **Scalable storage** host properties contain the following settings.

Table 2-54 Scalable storage host properties

Property	Description
Key Management Server (KMS) name	If you configured a key management service (KMS) server, the name of the primary server that sends the request to the KMS server is displayed here.
Metering interval	Determines how often NetBackup gathers connection information for reporting purposes. The value is set in seconds. The default setting is 300 seconds (5 minutes). If this value is set to zero, metering is disabled.
Total available bandwidth	Use this value to specify the speed of your connection to the cloud. The value is specified in kilobytes per second. The default value is 102400 KB/sec.
Sampling interval	The time, in seconds, between measurements of bandwidth usage. The larger this value, the less often NetBackup checks to determine the bandwidth in use. If this value is zero, throttling is disabled.
Advanced settings	Expand Advanced settings to configure additional settings for throttling. See "Configuring advanced bandwidth throttling settings" on page 157. See "Advanced bandwidth throttling settings" on page 157.
Maximum concurrent jobs	The default maximum number of concurrent jobs that the media server can run for the cloud storage server. This value applies to the media server, not to the cloud storage server. If you have more than one media server that can connect to the cloud storage server, each media server can have a different value. Therefore, to determine the total number of connections to the cloud storage server, add the values from each media server. If you configure NetBackup to allow more jobs than the number of connections, NetBackup fails any jobs that start after the number of maximum connections is reached. Jobs include both backup and restore jobs. You can configure job limits per backup policy and per storage unit. See "Limit jobs per policy (policy attribute)" on page 678. See "Maximum concurrent jobs storage unit setting" on page 563. Note: NetBackup must account for many factors when it starts jobs: the number of concurrent jobs, the number of connections per media server, the number of media servers, and the job load-balancing logic. Therefore, NetBackup may not fail jobs exactly at the maximum number of connections. NetBackup may fail a job when the connection number is slightly less than the maximum, exactly the maximum, or slightly more than the maximum. A value of 100 is generally not needed.

Configuring advanced bandwidth throttling settings

Advanced bandwidth throttling settings let you control various aspects of the connection between the NetBackup hosts and your cloud storage provider.

See [“Scalable Storage properties”](#) on page 155.

To configure advanced bandwidth throttling settings

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the media server.
- 4 If necessary, click **Connect**. Then click **Edit media server**.
- 5 Click **Scalable storage**.
- 6 Expand **Advanced settings**.
- 7 Configure the settings and then click **Save**.

See [“Advanced bandwidth throttling settings”](#) on page 157.

Advanced bandwidth throttling settings

The following table describes the advanced bandwidth throttling settings.

Table 2-55 Advanced throttling configuration settings

Property	Description
Read bandwidth	Use this field to specify the percentage of total bandwidth that read operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, restore or replication failures may occur due to time-outs. Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth. Default value: 100 Possible values: 0 to 100

Table 2-55 Advanced throttling configuration settings (*continued*)

Property	Description
Write bandwidth	Use this field to specify the percentage of total bandwidth that write operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, backup failures may occur due to time-outs. Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth. Default value: 100 Possible values: 0 to 100
Work time	Use this field to specify the time interval that is considered work time for the cloud connection. Specify a start time and end time. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Off time	Use this field to specify the time interval that is considered off time for the cloud connection. Specify a start time and end time. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Weekend	Specify the start and stop time for the weekend. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Read Bandwidth (KB/s)	This field displays how much of the available bandwidth the cloud storage server transmits to a NetBackup media server during each restore job. The value is expressed in kilobytes per second.

Table 2-55 Advanced throttling configuration settings (*continued*)

Property	Description
Write Bandwidth (KB/s)	This field displays how much of the available bandwidth the NetBackup media server transmits to the cloud storage server during backup jobs. The value is expressed in kilobytes per second.

Servers properties

To access this setting, in the NetBackup web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Servers**.

The **Servers** properties display the NetBackup server lists on the selected primary server, media server, or client. The server lists display the NetBackup servers that the host recognizes.

The **Primary server** field contains the name of the primary server for the selected host. (The name of the selected host appears in the title bar.)

The **Servers** page contains the following settings.

Table 2-56 Servers properties

Tab	Description
Additional servers tab	This tab lists the additional servers that can access the server that is specified as the Primary server. During installation, NetBackup sets the primary server to the name of the system where the server software is installed. NetBackup uses the primary server value to validate server access to the client. The primary server value is also used to determine which server the client must connect to so that files can be listed and restored. Note: For a Fibre Transport (FT) media server that has multiple network interfaces for VLANs: Ensure that the FT server's primary host name appears before any other interface names for that FT media server host. For more information, see the NetBackup SAN Client and Fibre Transport Guide.
Media servers tab	This tab lists the hosts that are media servers only. Hosts that are listed as media servers can back up and restore clients, but have limited administrative privileges. If you add a media server to both the Media servers tab and the Additional servers tab, this action may introduce unintended consequences. A computer that is defined as both a primary server and a media server gives the administrator of the media server full primary server privileges. You may inadvertently give the media server administrator more privileges than intended.

Table 2-56 Servers properties (*continued*)

Tab	Description
Trusted primary servers tab	Use this tab to add the remote primary servers that you trust using NetBackup CA-signed certificates or external CA-signed certificates and to view the primary servers that are already trusted. See “About trusted primary servers for Auto Image Replication” on page 974. See “Add a trusted primary server” on page 164. Note: If either the source or remote primary server is clustered, you must enable inter-node communication on all of the nodes in the cluster. Do so before you add the trusted primary server. See “Enabling NetBackup clustered primary server inter-node authentication” on page 161. Information about Auto Image Replication and storage lifecycle policies is available. If your user account is configured for multifactor authentication on the target host, append the one-time password to the password. See “About NetBackup Auto Image Replication” on page 961. See “About storage lifecycle policies” on page 601.

Adding a server to a servers list

Depending on the tab that is selected, you can add a primary server, media server, or client to the server list in the **Additional servers** tab or the **Media servers** tab.

To add a server to a servers list

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the host.
- 4 If necessary, click **Connect**. Then click **Edit primary server**, **Edit media server**, or **Edit client**.
- 5 Click **Servers**.
- 6 Select the tab that contains the server list that you want to modify.
- 7 Click **Add**.
- 8 Enter the name of the new server.
- 9 Click **Add**.

Note: If you add a media server, run `nbemmcmd -addhost` to add the media server to the Enterprise Media Manager (EMM) in the NetBackup database of the primary server.

Removing a server from a servers list

You can remove a primary server or a media server from the **Additional servers** list or the **Media servers** list.

To remove a server from a servers list

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the host.
- 4 If necessary, click **Connect**. Then click **Edit primary server**, **Edit media server**, or **Edit client**.
- 5 Click **Servers**.
- 6 Click the **Additional servers** tab or the **Media servers** tab.
- 7 Locate a server in the list.
- 8 Click **Actions > Delete**.

Enabling NetBackup clustered primary server inter-node authentication

NetBackup requires inter-node authentication among the primary servers in a cluster. For authentication, you must provision an authentication certificate on all of the nodes of the cluster. The certificates are used to establish SSL connections between the NetBackup hosts.

See [“Add a trusted primary server”](#) on page 164.

The inter-node authentication allows the following NetBackup functionality:

NetBackup web UI

The NetBackup web UI in primary server clusters requires the NetBackup authentication certificates for correct functionality.

Targeted A.I.R. (Auto Image Replication)

Auto Image Replication in which a primary server is in a cluster requires inter-node authentication among the hosts in that cluster. The NetBackup authentication certificates provide the means to establish the proper trust relationships.

Provision the certificates on the cluster hosts before you add the trusted primary server. This requirement applies regardless of whether the clustered primary server is the source of the replication operation or the target.

See [“About trusted primary servers for Auto Image Replication”](#) on page 974.

To enable inter-node authentication for a NetBackup clustered primary server

- ◆ On the active node of the NetBackup primary server cluster, run the following NetBackup command:

- **Windows:** `install_path\NetBackup\bin\admincmd\bpnbaz -setupat`
- **UNIX:** `/usr/opensv/netbackup/bin/admincmd/bpnbaz -setupat`

NetBackup creates the certificates on every node in the primary server cluster.

The following is example output:

```
# bpnbaz -setupat
You will have to restart Netbackup services on this machine after
the command completes successfully.
Do you want to continue(y/n)y
Gathering configuration information.
Please be patient as we wait for 10 sec for the security services
to start their operation.
Generating identity for host 'bitl.remote.example.com'
Setting up security on target host: bitl.remote.example.com
nbatd is successfully configured on Netbackup Primary Server.
Operation completed successfully.
```

About trusted primary servers

A trust relationship between NetBackup domains lets you do the following:

- Select specific domains as a target for replication. This type of Auto Image Replication is known as targeted A.I.R.

Without a trust relationship, NetBackup replicates to all defined target storage servers. A trust relationship is optional for Media Server Deduplication Pool and PureDisk Deduplication Pool as a target storage. To use a Cloud Catalyst storage server, a trust relationship is required.

- Include usage reporting for multiple primary servers.

Primary servers can use a NetBackup certificate authority (CA) certificate or an external CA certificate. NetBackup determines the CAs used by the source and the target domains and selects the appropriate CA to use for communication between the servers. If the target primary server is configured for both CA types, NetBackup prompts you to select the CA that you want to use.

About the certificate to use to add a trusted primary server

A source or a target primary server may use NetBackup CA-signed certificates (host ID-based certificates) or external CA-signed certificates.

For more information on NetBackup host ID-based certificates and external CA support, refer to the [NetBackup Security and Encryption Guide](#).

To establish trust between source and target primary servers, NetBackup verifies the following:

Can the source primary server establish trust using an external CA-signed certificate?	If the external CA configuration options - <code>ECA_CERT_PATH</code> , <code>ECA_PRIVATE_KEY_PATH</code> , and <code>ECA_TRUST_STORE_PATH</code> - are defined in the NetBackup configuration file of the source primary server, it can establish the trust using an external certificate. In the case of the Windows certificate trust store, only the option <code>ECA_CERT_PATH</code> is defined.
Which certificate authorities (CA) does the target primary server support?	The target primary server may support external CA, NetBackup CA, or both.

The following table lists the CA support scenarios and the certificate to use to establish trust between the source and the target primary servers. The instructions assume that you use the NetBackup web UI for the configuration.

Table 2-57 Certificate to use for the trust setup

Can the primary server use an external certificate?	Which CA does the target primary server use?	Certificate to use for the trust setup
Yes	External CA	External CA
The source primary server can use the NetBackup CA and an external CA for communication with a remote primary server.	NetBackup CA	NetBackup CA
	External CA and NetBackup CA	NetBackup prompts to select the CA that you want to use for trust setup.
No	External CA	No trust is established.
The source primary server can only use the NetBackup CA for communication with a remote primary server.	NetBackup CA	NetBackup CA
	External CA and NetBackup CA	NetBackup CA

Add a trusted primary server

Replication operations require that a trust relationship exists between the NetBackup servers in the different domains. You can create a trust relationship between the primary servers that both use the NetBackup CA or that both use an external CA.

Before you begin, review the following information:

- Ensure that you have the RBAC System Administrator role or a role with similar permissions. Or, for appliances with software versions 3.1 and later you must have permissions for the NetBackup CLI user.
- For a remote Windows primary server, the user's domain may not be the same as that of the authentication service. In this case you must add the domain with LDAP using the `vssat addldapdomain` command.
- For a NetBackup CA-signed certificate, the recommended method to authenticate the server is the option **Specify authentication token of the trusted primary server**.
- If you use the option **Specify credentials of the trusted primary server**, that method may present a possible security breach. Only an authentication token can provide restricted access and allow secure communication between both the hosts. To establish trust with a 3.1 NetBackup primary appliance, use the NetBackup CLI credentials.

To add a trusted primary server

- 1 Open the NetBackup web UI.
- 2 Identify the NetBackup versions that are installed and the certificate types that are used on the source and the target servers.

Both servers must use the same certificate type.
- 3 For the servers that use the NetBackup certificate authority (CA), obtain an authorization token for the remote server.
- 4 For the servers that use the NetBackup certificate authority (CA), obtain the fingerprint for each server.
- 5 At the top right, select **Settings > Global security**.
- 6 Select the **Trusted primary servers** tab.
- 7 Select the **Add** button.
- 8 Enter the fully-qualified host name of the remote primary server and select **Validate Certificate Authority**.
- 9 Follow the prompts in the wizard.
- 10 Repeat these steps on the remote primary server.

More information

For more information on using an external CA with NetBackup, see the [NetBackup Security and Encryption Guide](#).

Remove a trusted primary server

You can remove a trusted primary server, which removes the trust relationship between primary servers. Note the following implications:

- Any replication operations fail that require the trust relationship.
- A remote primary server is not included in any usage reporting after you remove the trust relationship.

To remove a trusted primary server, you must perform the following procedure on both the source and the target server.

To remove a trusted primary server

- 1 Open the NetBackup web UI.
- 2 Ensure that all replication jobs to the target primary server are complete.
- 3 Delete all storage lifecycle policies (SLPs) that use the trusted primary as a destination. Before deleting an SLP, ensure that there are no backup policies or protection plans that use the SLP for storage.

- 4 At the top right, select **Settings > Global security**.
- 5 Select the **Trusted primary servers** tab.
- 6 Locate the server that you want to remove.
- 7 Select **Actions > Remove**.
- 8 Select **Remove trust**.

Note: If you use multiple NICs, if you established trust using more than one host NIC and if you remove the trust relationship with any one host NIC, the trust with all the other host NICs is broken.

Changing the primary server that performs backups and restores for a client

Use the **Make primary** option to change the primary server that performs backups and restores for a client. This option does not change a host into a primary server.

Note: The client can also change their primary server in the **Backup, Archive, and Restore** interface by selecting **Actions > Specify NetBackup Machines and Policy Type**. In this dialog, select the primary server to use for backups and restores.

This option is useful in a disaster recovery situation or in a NetBackup environment where Auto Image Replication is configured. For example, select a client in the source domain, then use the **Make primary** option to temporarily point the client to the primary server of the target domain. After you change the primary server, restores from the target domain can be initiated.

To change the primary server that a client uses for backups and restores

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the client.
- 4 If necessary, click **Connect**. Then click **Edit client**.
- 5 Click **Servers**.

6 On the **Additional servers** tab, locate the server.

7 Click **Actions > Make primary**.

In the configuration file, the new primary server appears as the first server entry in the list.

Changing the primary server does not prevent the former primary server from initiating backups for the client. As long as that server continues to be listed on the client's server list, the primary server can perform backups.

SharePoint properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **SharePoint**.

The **SharePoint** properties protect SharePoint Server installations and apply to the currently selected Windows client.

For complete information on these options, see the [NetBackup for Microsoft SharePoint Server Administrator's Guide](#).

The **SharePoint** host properties contain the following settings.

Table 2-58 SharePoint host properties

Property	Description
Domain\Username	Specifies the domain and the user name for the account you want to use to log on to SharePoint (DOMAIN\user name).
Password	Specifies the password for the account.
Consistency check before backup	Specifies the consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation. These checks are performed for both server-directed and user-directed backups. If you choose to perform a consistency check, you can select Continue with backup if consistency check fails. NetBackup then continues to perform the backup if the consistency check fails.
SharePoint granular restore proxy host	For any VMware backups that protect Federated SharePoint configurations, provide the name of the back-end SQL server. This server acts as the granular restore proxy host for the catalog hosts (front-end servers in the farm).

Consistency check options for SharePoint Server

The following consistency checks can be performed before a SharePoint Server backup.

Table 2-59 Consistency check options

Option	Description
None	Do not perform consistency checking.
Full check, excluding indexes	Select this option to exclude indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the non-clustered index pages is not checked.
Full check, including indexes	Include indexes in the consistency check. Any errors are logged.

SLP settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **SLP settings**. You can also configure the SLP settings from **Storage > Storage lifecycle policies > SLP settings**.

The **SLP settings** properties allow administrators to customize how storage lifecycle policies (SLPs) are maintained and how SLP jobs run. These properties apply to the SLPs of the currently selected primary server.

[Table 2-60](#) describes the available properties for SLPs. It also lists the syntax to use with the command-line method.

Use the list in the **Units** column to change the units of measurement for the size or the time.

Table 2-60 SLP settings

Property	Description
Minimum size per duplication job	The smallest batch size that can run as a single duplication job. The job does not run until enough images accumulate to reach this minimum batch size or until the Force interval for small jobs time is reached. Minimum: 1 kilobyte; no maximum size. Default: 8 gigabytes. Configuration option default: <code>SLP.MIN_SIZE_PER_DUPLICATION_JOB = 8 GB</code>

Table 2-60 SLP settings (*continued*)

Property	Description
Maximum size per duplication job	The largest batch size that can run as a single duplication job. Minimum: 1 kilobyte; no maximum size. Default: 100 gigabytes. Configuration entry default: <code>SLP.MAX_SIZE_PER_DUPLICATION_JOB = 100 GB</code>
Maximum size per A.I.R. replication job	The largest batch size that can run as a single job for Auto Image Replication. Minimum: 1 kilobyte; no maximum size. Default: 100 gigabytes. Configuration entry default: <code>SLP.MAX_SIZE_PER_BACKUP_REPLICATION_JOB = 100 GB</code>
Maximum images per snapshot replication job	The largest number of images in a single batch that can run as a single job. Default: 50 images, with no minimum number or maximum number. Use this parameter with the Limit I/O streams disk pool option which limits the number of jobs that can run concurrently to each volume in the disk pool. Configuration entry default: <code>SLP.MAX_IMAGES_PER_SNAPSHOT_REPLICATION_JOB = 50</code>
Minimum images per A.I.R. Import job	The fewest number of images in a single batch that can run as an Auto Image Replication import job. The job does not run until either the minimum size is reached or the Force interval for small jobs time is reached. Minimum: 1 image; no maximum number of images. Default: 1 image. Configuration entry default: <code>SLP.MIN_IMAGES_PER_IMPORT_JOB = 1</code>
Maximum images per A.I.R. Import job	The largest number of images in a single batch that can run as an Auto Image Replication import job. Minimum: 1 job; no maximum number of images. Default: 250 images. Configuration entry default: <code>SLP.MAX_IMAGES_PER_IMPORT_JOB = 250</code>
Force interval for small jobs	The age that the oldest image in a batch must reach after which the batch is submitted as a duplication job. This value prevents many small duplication jobs from running at one time or running too frequently. It also prevents NetBackup from waiting too long before it submits a small job. Default: 30 minutes, with no minimum number or maximum number. Configuration entry default: <code>SLP.MAX_TIME_TIL_FORCE_SMALL_DUPLICATION_JOB = 30 MINUTES</code>

Table 2-60 SLP settings (*continued*)

Property	Description
Job submission interval	Indicates the frequency of the job submission for all operations. No minimum interval or maximum interval. Default: 5 minutes. By default, all jobs are processed before more jobs are submitted. Increase this interval to allow NetBackup to submit more jobs before all jobs are processed. Set the interval when the list of available images is scanned for those that can be batched together and jobs submitted. A shorter interval allows for a better response to changing system workloads at the cost of increased processing. Configuration entry default: <code>SLP.JOB_SUBMISSION_INTERVAL = 5 MINUTES</code>
Image processing interval	The number of minutes between image-processing sessions. Set the interval when newly created images are recognized and set up for SLP processing. Default: 5 minutes. Configuration entry default: <code>SLP.IMAGE_PROCESSING_INTERVAL = 5 MINUTES</code>
Cleanup interval	The time between when a job finishes and before NetBackup removes the job artifacts for the completed job. No minimum interval or maximum interval. Default: 24 hours. Configuration entry default: <code>SLP.CLEANUP_SESSION_INTERVAL = 24 HOURS</code>
Extended image retry interval	The amount of time to wait before an unsuccessful operation is added to the first job that runs after the delay. (This behavior applies to all SLP jobs.) The extra time gives the administrator additional time to solve a problem that prevents job completion. No minimum interval or maximum interval. Default: 2 hours. Configuration entry default: <code>SLP.IMAGE_EXTENDED_RETRY_PERIOD = 2 HOURS</code>
Unused SLP definition version cleanup delay	Concerns the deletion of SLP versions where a more recent version exists. The setting controls how long a version must be inactive before NetBackup deletes it. Default: 14 days. Configuration entry default: <code>SLP.VERSION_CLEANUP_DELAY = 14 DAYS</code> See "Deleting old storage lifecycle policy versions" on page 645.

Table 2-60 SLP settings (*continued*)

Property	Description
Tape resource multiplier	Limits the number of concurrently active duplication jobs that can access a single tape media storage unit to xx times the number of available drives. Allows tuning to avoid overloading the Resource Broker, yet makes sure that the devices are not idle. No minimum multiplier or maximum multiplier. Default: 2 (multiply access to the write drives by two). Configuration entry default: <code>SLP.TAPE_RESOURCE_MULTIPLIER = 2</code>
Disk resource multiplier	Limits the number of concurrently active duplication jobs that can access a single disk storage unit to xx times the number of available drives. Allows tuning to avoid overloading the Resource Broker, yet makes sure that the devices are not idle. No minimum multiplier or maximum multiplier. Default: 2 (multiply access to the write drives by two). Configuration entry default: <code>SLP.DISK_RESOURCE_MULTIPLIER = 2</code>
Group images across SLPs	If this parameter is set to Yes (default), multiple SLPs of the same priority can be processed in the same job. If No, batching can occur only within a single SLP. Configuration entry default: <code>SLP.DUPLICATION_GROUP_CRITERIA = 1</code> Configuration entry for no, do not allow batching: <code>SLP.DUPLICATION_GROUP_CRITERIA = 0</code>
Window close buffer time	Sets the amount of time before a window closes when NetBackup does not submit new jobs using that window. Minimum 2 minutes; maximum: 60 minutes. Default: 15 minutes. Configuration entry default: <code>SLP.WINDOW_CLOSE_BUFFER_TIME = 15 MINUTES</code>
Deferred duplication offset time	For deferred operations, jobs are submitted x time before the source copy is due to expire. Default: 4 hours. Configuration entry default: <code>SLP.DEFERRED_DUPLICATION_OFFSET_TIME = 4 HOURS</code>
Auto create A.I.R. Import SLP	Used for Auto Image Replication, indicates whether an SLP (that contains an Import operation) is created automatically in the target domain if no SLP is configured there. Default: Yes, an SLP is created in the target domain. Configuration entry default: <code>SLP.AUTO_CREATE_IMPORT_SLP = 1</code>

Table 2-60 SLP settings (*continued*)

Property	Description
How long to retry failed A.I.R. import jobs	How long NetBackup retries an Import job before it stops and deletes the record. After the initial four attempts, the retries become less frequent. Default: 0 (do not retry after the initial four attempts). Configuration entry default: <code>SLP.REPLICA_METADATA_CLEANUP_TIMER = 0 HOURS</code>
Pending A.I.R import threshold	How long NetBackup waits before it generates a notification that an Auto Image Replication copy is still in import pending state. After an Auto Image Replication copy has been replicated, NetBackup puts the source copy into import pending state. If the copy is in import pending state for the time period that this threshold sets, NetBackup generates a notification. Notifications are sent to the NetBackup error log and are visible in the Problems report. Notifications may also be sent to an email address, if specified. Default: 24 hours Configuration entry default: <code>SLP.PENDING_IMPORT_THRESHOLD = 24 HOURS</code> See “About Auto Image Replication import confirmation” on page 981.
Email address to receive notifications	The email address that receives pending A.I.R. import notifications. Default: None. Configuration entry format: <code>SLP.NOTIFICATIONS_ADDRESS = user@company.com</code>

Using the command line to change SLP parameters

You can also change the parameters using the command line.

To use the command-line method, use the `nbgetconfig` and the `nbsetconfig` commands to change the defaults. For information about these commands, see the [NetBackup Commands Reference Guide](#).

Command-line units of measurement for the SLP parameters

The abbreviations are case-insensitive for units of measurement.

The following abbreviations can be used where sizes are indicated:

bytes	kb	kilobyte	kilobyte(s)	kilobytes	mb	megabyte
megabyte(s)	megabytes	gb	gigabyte	gigabyte(s)	gigabytes	tb
terabyte	terabyte(s)	terabytes	pb	petabyte	petabyte(s)	petabytes

The following abbreviations can be used where units of time are indicated:

sec	second	second(s)	seconds	min	minute	minute(s)	minutes
hour	hour(s)	hours	day	day(s)	days	mon	month
month(s)	months	week	week(s)	weeks	year	year(s)	years

`nbcl.conf` file

Whenever a storage lifecycle policy parameter is changed from the default, the change creates the `nbcl.conf` configuration file.

This file is found in the following locations. It is present only if the default of any parameter has been changed.

- On Windows:

`install_path\NetBackup\var\global\nbcl.conf`

- On UNIX:

`/usr/opensv/var/global/nbcl.conf`

About batch creation logic in Storage Lifecycle Manager

The Storage Lifecycle Manager service (`nbstserv`) is in charge of creating duplication jobs for storage lifecycle policies. Part of duplication job creation includes grouping the backup (or source) jobs into batches.

Note: Restart `nbstserv` after making changes to the underlying storage for any operation in an SLP.

One objective of the batching logic is to prevent media contention for tape operations, including virtual tape libraries (VTL).

Batching logic applies to both disk and tape. (Though the method to prevent media contention for disk is to use disk pools and then to limit I/O streams to disk pools.)

The batching logic requires that for each evaluation cycle, `nbstserv` consider all completed source jobs when determining which duplication job to run next. By default, `nbstserv` performs the evaluation once every 5 minutes.

`nbstserv` avoids overloading the Resource Broker (`nbrb`) queue with jobs. Too many jobs in the queue make the role of the Resource Broker harder and slows down system performance.

By default, `nbstserv` now creates groups based on the **Group images across SLPs** parameter in the **SLP Parameters** host properties. By default, multiple storage lifecycle policies with the same priority can be batched together.

See “[SLP settings properties](#)” on page 168.

This batching logic change affects how duplication jobs appear in the **Activity Monitor**. Storage lifecycle policies that have been combined into one job appear under a single policy name: `SLP_MultipleLifecycles`. If a storage lifecycle policy has not been combined with another, the name appears in the **Activity Monitor** under the name of the SLP: `SLP_name`.

Users may see some duplication jobs that, although in the running state, do not duplicate data because they have no resources to read or write. These jobs continue to run until they receive resources to complete the job.

To turn off grouping by duplication job priority, set the **Group images across SLPs** parameter to **No** in the **SLP Parameters** host properties.

Throttle bandwidth properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Throttle bandwidth**.

Use the **Throttle bandwidth** properties to specify a limit for the network bandwidth or transfer rate that NetBackup clients use on a network. The actual limiting occurs on the client side of the backup connection. These properties limit only backups. Restores are unaffected. The default is that the bandwidth is not limited.

The **Throttle bandwidth** properties are similar to the **Bandwidth** host properties, but offer greater flexibility in IPv6 environments.

To add, edit, or remove a throttle bandwidth setting

- 1 Open the **NetBackup web UI**.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Throttle bandwidth**.

Add a setting	To add a network or host setting	
	1	Click Add .
	2	Enter the name of the network or host to which the throttle applies.
	3	Select the bandwidth for the network or host indicated. A value of zero disables the throttling of IPv6 addresses. This value is the transfer rate in kilobytes per second. A value of zero disables the throttling of IPv6 addresses.
	4	Click Add .
Edit a setting	To edit a network or host setting	
	1	Locate the name of the network or host.
	2	Click Actions > Edit .
	3	Make the wanted changes.
	4	Click Save .
Delete a setting	Delete a a network or host setting	
	1	Locate the name of the network or host.
	2	Click Actions > Delete .

5 Click **Save**

See [“Bandwidth properties”](#) on page 55.

Timeouts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Timeouts**.

The **Timeouts** properties apply to the selected primary server, media server, or client.

Table 2-61 Timeouts host properties

Property	Description
Client connect timeout	This property applies to the currently selected server. Specifies the number of seconds the server waits before it times out when it connects to a client. The default is 300 seconds.

Table 2-61 Timeouts host properties (*continued*)

Property	Description
Backup start notify timeout	This property applies to the currently selected server . Specifies the number of seconds the server waits for the <code>bpstart_notify</code> script on a client to complete. The default is 300 seconds. Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can time out while the <code>bpstart_notify</code> script is running.
Media server connect timeout	This property applies to the currently selected server . Specifies the number of seconds that the primary server waits before it times out when it connects to a remote media server. The default is 30 seconds.
Client read timeout	This property applies to the currently selected server or client. Specifies the number of seconds that NetBackup waits for a response from a client before the operation attempt fails. This timeout can apply to a NetBackup primary, remote media server, or database-extension client (such as NetBackup for Oracle). The default is 300 seconds. If the server does not get a response from a client within the Client read timeout period, the backup or the restore operation can fail. See the section called “Recommendations for the Client read timeout” on page 177. The sequence on a database-extension client is as follows: ■ NetBackup on the database-extension client reads the client’s client-read timeout to find the initial value. If the option is not set, the standard 5-minute default is used. ■ When the database-extension API receives the server’s value, it uses it as the client-read timeout.
Backup end notify timeout	This property applies to the currently selected server. Specifies the number of seconds that the server waits for the <code>bpend_notify</code> script on a client to complete. The default is 300 seconds. Note: If this timeout is changed, verify that Client read timeout is set to the same or higher value.

Table 2-61 Timeouts host properties (*continued*)

Property	Description
Use OS dependent timeouts	This property applies to the currently selected server or client. Specifies that the client waits for the timeout period as determined by the operating system when it lists files, as follows: ■ Windows client: 300 seconds ■ UNIX client: 1800 seconds File browse timeout Specifies how long the client can wait for a response from the NetBackup primary server while it lists files. If the limit is exceeded, the user receives a socket read failed error. The timeout can be exceeded even while the server processes the request. Note: If it exists, the value in a UNIX client's <code>\$HOME/bp.conf</code> file takes precedence to the property here.
Media mount timeout	This property applies to the currently selected primary server. Specifies how long NetBackup waits for the requested media to be mounted, positioned, and ready on backups, restores, and duplications. Use this timeout to eliminate excessive waiting time during manual media mounts. (For example, when robotic media is out of the robot or is off-site.)

Recommendations for the Client read timeout

It is recommended to increase the timeout value in the following situations:

- The client-read timeout on a database-extension client is a special case. Clients can initially require more time to get ready than other clients. More time is required because database backup utilities frequently start several backup jobs at the same time, slowing the central processing unit. A setting of 15 minutes is adequate for many installations.
- Backing up directly to an MSDP cloud storage server. If the value is not increased for both the primary server and the media server, you may see jobs failing with the following message in the job details:

```
Error bpbrm (pid=119850) socket read failed: errno = 62 - Timer expired
```

Note that increasing the timeout is not needed if you use a storage lifecycle policy to first back up to an MSDP storage server and then duplicate the data to an MSDP cloud storage server using an optimized duplication operation. (This operation is the recommended method of operation.)

Note: If using the `bpstart_notify` script: The **Client read timeout** (`CLIENT_READ_TIMEOUT` option) must be equal to or greater than the **Backup start notify timeout** (`BPSTART_TIMEOUT` option). If the **Client read timeout** is less than the **Backup start notify timeout**, the job can timeout while the `bpstart_notify` script is running.

Universal settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Universal settings**.

Use the **Universal settings** properties to configure certain backup and restore settings. These properties apply to a selected primary server, media server, or client.

The **Universal settings** host properties contain the following settings.

Table 2-62 Universal settings properties

Property	Description
Restore retries	This setting applies to the selected server or client. Specifies the number of attempts a client has to restore after a failure. (The default is 0; the client does not attempt to retry a restore. The client can try up to three times.) Change Restore retries only if problems are encountered. If a job fails after the maximum number of retries, the job goes into an incomplete state. The job remains in the incomplete state as determined by the Move restore job from incomplete state to done state property. See “Clean up properties” on page 61. A checkpointed job is retried from the start of the last checkpointed file rather than at the beginning of the job. Checkpoint restart for restore jobs allows a NetBackup administrator to resume a failed restore job from the Activity Monitor. See “Take checkpoints every __ minutes (policy attribute)” on page 674.

Table 2-62 Universal settings properties (*continued*)

Property	Description
Browse timeframe for restores	This setting applies to the selected server and applies to all NetBackup clients. Specifies the timeframe that NetBackup uses to search for files to restore. By default, NetBackup includes files from the time of the last-full backup through the latest backup for the client. ■ Timeframe. Specifies how long ago NetBackup searches for files to restore. For example, to limit the browse range to one week before the current date, select Timeframe and specify 7. ■ Last full backup. Indicates whether NetBackup includes all backups since the last successful full backup in its browse range. This option is enabled by default. If the client belongs to more than one policy, then the browse starts with the earliest of the set of last-full backups.
Use specified network interface	This setting applies to the selected server or client. Specifies the network interface that NetBackup uses to connect to another NetBackup client or server. A NetBackup client or server can have more than one network interface. To force NetBackup connections to be made on a specific network interface, use this entry to specify the network host name of that interface. By default, the operating system determines the one to use.
Allow server file writes	This setting applies to the selected server or client. Specifies whether a NetBackup server can create or modify files on the NetBackup client. For example, disable this property to prevent server-directed restores and remote changes to the client properties. After the Allow server file writes property is applied, it can be cleared only by modifying the client configuration. The default is that server writes are allowed.
Administrator	This setting applies to the selected server or client. Specifies whether the server or the client sends email. ■ Server sends mail With this option the server sends an email to the address that is specified in the Global attributes properties. Enable this property if the client cannot send mail and you want an email notification. The default is that this property is disabled. See “Global attributes properties” on page 109. ■ Client sends mail With this option the client sends an email to the address that is specified in the Universal settings properties. If the client cannot send email, use Server sends mail. The default is that this property is enabled.

Table 2-62 Universal settings properties (*continued*)

Property	Description
Client administrator's email	Specifies the email address of the administrator on the client. This address is where NetBackup sends backup status reports for the client. By default, no email is sent. To enter multiple addresses or email aliases, separate entries with commas.

User account settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **User account settings**.

Use the **User account settings** properties to customize the settings for user sessions, user account lockout, and the sign-in banner.

Table 2-63 User account settings properties

Property	Description
Session idle timeout	Logs out the user session if there is no activity for the specified period of time. See "Configure when idle sessions should time out" on page 182.
Maximum concurrent sessions	Limits the number of sessions that a user can have open concurrently. See "Configure the maximum of concurrent user sessions" on page 182.
User account lockout	Lock out an account after the specified number of failed sign-in attempts. See "Configure the maximum of failed sign-in attempts" on page 183.
Sign-in banner configuration	You can configure a sign-in banner that displays each time that any user signs in to the NetBackup web UI. A different banner can be configured for any primary server. See "Display a banner to users when they sign in" on page 183.

Terminate a NetBackup user session

For security or maintenance purposes, you can terminate one or more NetBackup user sessions. To configure NetBackup to automatically terminate any idle user sessions, see the following topic.

See ["Configure when idle sessions should time out"](#) on page 182.

Note: Changes to a user's roles are not immediately reflected in the web UI. An administrator must terminate the active user session before any changes take effect. Or, the user must sign out and sign in again.

To sign out a user session

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Go to the **Active sessions** tab.
- 5 Select the user session that you want to sign out.
- 6 Select **Terminate session**.

To sign out all user sessions

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Go to the **Active sessions** tab.
- 5 Select **Terminate all sessions**.

Unlock a NetBackup user

You can view the user accounts that are currently locked out of NetBackup and unlock one or more users.

By default a user's account only remains locked for 24 hours. You can change this time by adjusting the **User sessions > User account settings > User account lockout** setting.

See [“Configure the maximum of failed sign-in attempts”](#) on page 183.

To unlock out a locked user account

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Go to the **Locked users** tab.

- 5 Select the user account that you want to unlock.
- 6 Select **Unlock**.

To unlock all locked user accounts

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Go to the **Locked users** tab.
- 5 Select **Unlock all users**.

Configure when idle sessions should time out

You can customize when user sessions should time out and a user is automatically signed out. The setting you choose is applied to the NetBackup web UI. To configure this setting from the command line, use `nbsetconfig` to set the `GUI_IDLE_TIMEOUT` option.

To configure when idle sessions should time out

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Turn on **Session idle timeout** and click **Edit**.
- 5 Select the number of minutes and click **Save**.

For active users, the updates are applied the next time the user signs in.

Configure the maximum of concurrent user sessions

This setting limits the number of concurrent API sessions that a user can have active. This setting does not apply to API key sessions or to other applications like the NetBackup Backup, Archive, and Restore interface.

To configure this setting from the command line, use `nbsetconfig` to set the `GUI_MAX_CONCURRENT_SESSIONS` option.

To configure the maximum of concurrent user sessions

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.

- 4 Turn on **Maximum concurrent sessions** and click **Edit**.
 - 5 Select the **Number of concurrent sessions per user** and click **Save**.
- For active users, the updates are applied the next time the user signs in.

Configure the maximum of failed sign-in attempts

You can automatically lock a user account if the user exceeds a maximum number of failed sign-in attempts. The user account remains locked until the account lockout period passes.

If there is an immediate need to access NetBackup, the administrator can unlock the account.

See [“Unlock a NetBackup user”](#) on page 181.

You can customize the maximum number of NetBackup failed sign-in attempts. The setting you choose applies only to the NetBackup web UI. To configure this setting from the command line, use `nbsetconfig` to set the `GUI_MAX_LOGIN_ATTEMPTS` and `GUI_ACCOUNT_LOCKOUT_DURATION` options.

To configure the maximum of failed sign-in attempts

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Turn on **User account lockout** and click **Edit**.
- 5 Select the number of failed sign-in attempts that you want to allow before an account is locked.
- 6 To unlock a locked account after a period of time, select the number of minutes for **Unlock locked accounts after**.
- 7 Select **Save**.

For active users, the updates are applied the next time the user signs in.

Display a banner to users when they sign in

You can configure a sign-in banner that displays each time that any user signs in to the NetBackup web UI. A different banner can be configured for any primary server. This banner can also require the user to agree to the terms of service before the user signs in.

To display a banner to users when they sign in

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Turn on **Sign-in banner configuration** and click **Edit**.
- 5 Enter the text you want to use for the heading and the body of the message.
- 6 If you want to require the user to agree to the terms of service, select **Include "Agree" and "Disagree" buttons on the sign-in banner**.
- 7 Select **Save**.

For active users, the updates are applied the next time the user signs in.

To remove the sign-in banner

- 1 Open the web UI.
- 2 On the left, click **Security > User sessions**.
- 3 At the top right, click **User account settings**.
- 4 Turn off **Sign-in banner configuration**
- 5 Select **Save**.

For active users, the updates are applied the next time the user signs in.

UNIX client properties

Use the **UNIX client** properties to define properties of clients running on the UNIX platform.

See [“Busy file settings properties”](#) on page 59.

See [“Client settings properties for UNIX clients”](#) on page 73.

See [“Lotus Notes properties”](#) on page 117.

UNIX Server properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Linux primary server. If necessary click **Connect**, then and click **Edit primary server**. Then click **UNIX server**.

Use the **UNIX server** properties to change the **NFS access timeout** property. This property specifies how long the backup waits to process the mount table before it considers an NFS file system unavailable. The default is 5 seconds.

These properties apply to selected Linux primary servers.

See [“Follow NFS \(policy attribute\)”](#) on page 682.

VMware access hosts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **VMware access hosts**.

You can also access this setting from **Workloads > VMware > VMware settings > Access hosts**.

Use the **VMware access hosts** host properties to add or remove VMware backup hosts. These properties apply to currently selected primary servers.

These properties appear when the NetBackup Enterprise Client license is installed.

The backup host is a NetBackup client that performs backups on behalf of the virtual machines. (This host was formerly known as the VMware backup proxy server.) The backup host is the only host on which NetBackup client software is installed. As an option, the backup host can also be configured as a NetBackup primary server or media server.

The backup host is referred to as the recovery host when it performs a restore

You can add servers to and remove servers from the access hosts list:

Add Click **Add** and enter the fully qualified domain name of the backup host.

Remove Locate the backup host in the list and click **Remove**.

For more information, see the [NetBackup for VMware Administrator's Guide](#).

Windows client properties

Use the **Windows client** properties to configure specific NetBackup properties for Windows clients.

See [“Client settings properties for Windows clients”](#) on page 77.

See [“Lotus Notes properties”](#) on page 117.

See [“Exchange properties”](#) on page 92.

See [“SharePoint properties”](#) on page 167.

See [“Active Directory properties”](#) on page 55.

See [“Enterprise Vault properties”](#) on page 90.

Configuration options not found in the host properties

Most NetBackup configuration options can be found in the **Host properties** of the **NetBackup web UI**. However, some options cannot be accessed in the **Host properties**.

To change the default value for an option that is not found in the **Host properties**, first use the `nbgetconfig` command to obtain a list of configuration options. Then use `nbsetconfig` to change the options as needed.

For information about these commands, see the following resources:

- [NetBackup Commands Reference Guide](#)
- See [“About using commands to change the configuration options on UNIX or Linux clients and servers”](#) on page 186.

For information about the configuration options that are available, see the following topics:

- See [“Configuration options for NetBackup servers”](#) on page 187.
- See [“Configuration options for NetBackup clients”](#) on page 329.

About using commands to change the configuration options on UNIX or Linux clients and servers

When commands (`nbsetconfig` or `bpsetconfig`) are used to change the configuration options on UNIX or Linux NetBackup servers or clients, the commands change the appropriate configuration files.

Most options are found in the following configuration file:

```
/usr/opensv/netbackup/bp.conf
```

If a single UNIX or Linux system is running as both a client and a server, the `bp.conf` file contains options for both the client and the server.

The `bp.conf` file observes the following syntax:

- Use the # symbol to comment out lines.
- Any number of spaces or tabs are allowed on either side of = signs.
- Blank lines are allowed.
- Any number of blanks or tabs are allowed at the start of a line.

Each nonroot user on a UNIX or Linux client can also have a personal `bp.conf` file in their home directory:

```
$HOME/bp.conf
```

The options in personal `bp.conf` files apply only to user operations. During a user operation, NetBackup checks the `$HOME/bp.conf` file before `/usr/opensv/netbackup/bp.conf`.

Root users do not have personal `bp.conf` files. NetBackup uses the `/usr/opensv/netbackup/bp.conf` file for root users.

Stop and restart all NetBackup daemons and utilities on the server after you make a change to the `bp.conf` file on a Linux primary server. This action ensures that all of the NetBackup processes use the new `bp.conf` values. This action is not required for changes to `bp.conf` files on a client or to a `$HOME/bp.conf` file on the primary server.

The `SERVER` option must be present in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX or Linux clients and servers. During installation, NetBackup sets the `SERVER` option to the name of the primary server where the software is installed. It is the only required option in the `bp.conf` files. NetBackup uses internal software defaults for all options in the `bp.conf` file, except `SERVER`.

The `SERVER` entries must be the same on all servers in a primary and a media server cluster. It is recommended that all other entries also match on all servers. (The `CLIENT_NAME` option is an exception.)

Configuration options for NetBackup servers

The following topics are about configuration options for NetBackup servers. Nearly all of these options can also be set in the Host properties in the **NetBackup web UI**.

Note: On Windows platform, NetBackup supports 7-bit ASCII characters for the file paths that are to be specified for security-specific configuration options.

ALLOW_MEDIA_OVERWRITE option for NetBackup servers

This option overrides the NetBackup overwrite protection for various media formats on removable media.

Table 2-64 ALLOW_MEDIA_OVERWRITE information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ALLOW_MEDIA_OVERWRITE = media_format</pre> This option should appear only once in the configuration file.
Example	On the primary server (and media servers if applicable), add the following entry to permit overwriting the <code>cpio</code> format: <pre>ALLOW_MEDIA_OVERWRITE = CPIO</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Media > Allow media overwrite. See "Media properties" on page 119.

AUTO_ADD_ALL_ALIASES_FOR_CLIENT option for NetBackup servers

This option allows client aliases to be automatically added to the NetBackup database when `bpdbm` detects a new client in a backup policy.

Table 2-65 AUTO_ADD_ALL_ALIASES_FOR_CLIENT information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-65 **AUTO_ADD_ALL_ALIASES_FOR_CLIENT** information
(continued)

Usage	Description
How to use	By default, <code>AUTO_ADD_ALL_ALIASES_FOR_CLIENT</code> is not present in the configuration file. When <code>AUTO_ADD_ALL_ALIASES_FOR_CLIENT</code> is not present, the option is enabled. That is, <code>bpdbm</code> is allowed to add client aliases automatically. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>AUTO_ADD_ALL_ALIASES_FOR_CLIENT = YES NO</pre> This entry should appear only once in the configuration file.
Example	The following entry prohibits <code>bpdbm</code> from adding a client alias automatically: <pre>AUTO_ADD_ALL_ALIASES_FOR_CLIENT = NO</pre>
Equivalent NetBackup web UI property	No equivalent exists in the NetBackup web UI host properties.

BPBRM_VERBOSE option for NetBackup servers

The `BPBRM_VERBOSE` option is used for debugging purposes. It controls the amount of information that NetBackup includes in the `bpbrm` debug log.

Table 2-66 **BPBRM_VERBOSE** information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-66 BPBRM_VERBOSE information (continued)

Usage	Description
How to use	The default is that <code>BPBRM_VERBOSE</code> is the same value as the <code>VERBOSE</code> option (Global logging level). The <code>BPBRM_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPBRM_VERBOSE = -1 0 1 2 3 4 5</pre> This entry should appear only once in the configuration file.
Example	■ To use the same value as the <code>VERBOSE</code> option, enter: <code>BPBRM_VERBOSE = 0</code> This option is the same as setting the BPBRM logging level in the Logging host properties to Same as Global. ■ To log the minimum amount of information, enter: <code>BPBRM_VERBOSE = -1</code> This option is the same as setting the BPBRM logging level to 0 in the Logging host properties. ■ To log additional information, enter a value of 1 through 5: <code>BPBRM_VERBOSE = 1</code> This option is the same as setting the BPBRM logging level to 1 in the Logging host properties. ■ To log the maximum amount of information, enter: <code>BPBRM_VERBOSE = 5</code> This option is the same as setting the BPBRM logging level to 5 in the Logging host properties.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Logging > BPBRM logging level. See the NetBackup Logging Reference Guide for more information about the debug log. See “Logging properties” on page 113.

BPCD_ALLOWED_PATH option for NetBackup servers and clients

NetBackup denies access to a file that is specified for NetBackup operations if the path is a non-default path. For example, a path that is specified for progress log or rename files.

You must use the `BPCD_ALLOWED_PATH` option to allow access to non-default custom paths.

Table 2-67 BPCD_ALLOWED_PATH information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPCD_ALLOWED_PATH = Absolute pathname to the directory</pre>
Example	The following are the example entries on a NetBackup server or client: <pre>BPCD_ALLOWED_PATH = directory1</pre> <pre>BPCD_ALLOWED_PATH = directory2</pre> Create a separate entry for each directory.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

BPDBJOBS_COLDEFS options for Linux primary servers

Use `BPDBJOBS_COLDEFS` entries to customize the output of the `bpdjobs` process. Add a `BPDBJOBS_COLDEFS` option for every column you want to include in the output.

Add `BPDBJOBS_COLDEFS` entries to the `bp.conf` file to customize the output of the `bpdjobs` process.

Table 2-68 BPDBJOBS_COLDEFS information

Usage	Description
Where to use	On a Linux NetBackup primary server.

Table 2-68 BPDBJOBS_COLDEFS information (*continued*)

Usage	Description
How to use	Add BPDBJOBS_COLDEFS to the <code>/usr/opensv/netbackup/bp.conf</code> file. Add an entry for every column to include in the output by using the following format: <pre>BPDBJOBS_COLDEFS = COLDEFS_ENTRY [minimum_size [true false]]</pre> The following variables are defined: ■ COLDEFS_ENTRY is the name of the column to include in the output. ■ minimum_size is the minimum column width. If not specified, the default is a width of 5. ■ true indicates that the column should expand as needed. If not specified, true is the default. ■ false indicates that the column should not expand beyond the minimum_size.
Example	The order of the entries determines the order in which the column headings appear. <pre>BPDBJOBS_COLDEFS = JOBID 5 true BPDBJOBS_COLDEFS = TYPE 4 true BPDBJOBS_COLDEFS = STATE 5 true BPDBJOBS_COLDEFS = STATUS 6 true BPDBJOBS_COLDEFS = POLICY 6 true BPDBJOBS_COLDEFS = SCHEDULE 8 true BPDBJOBS_COLDEFS = CLIENT 6 true BPDBJOBS_COLDEFS = DSTMEDIA_SERVER 12 true BPDBJOBS_COLDEFS = ACTPID 10 true</pre> The appearance of BPDBJOBS_COLDEFS entries in the <code>bp.conf</code> file has the following ramifications: ■ The addition of any BPDBJOBS_COLDEFS option overrides all default columns. ■ All users on the local system see only those columns that are specified in the <code>bp.conf</code> file.
Equivalent host property	No equivalent exists in the host properties.

Table 2-69 shows possible COLDEFS entries and the column which is created by each.

Table 2-69 COLDEFS entries

COLDEFS entry	Column Name
ACTIVEELAPSED	Active Elapsed (elapsed active time)

Table 2-69 COLDEFS entries (*continued*)

COLDEFS entry	Column Name
ACTPID	Active PID (PID of job)
ATTEMPT	Attempt
BACKUPTYPE	Backup Type
CLIENT	Client
COMPLETION	Completion (percent complete)
COMPRESSION	Compression (yes or no)
COMPRESSION_SPACE_RATIO	Compression Space
DEDUPRATIO	Dedupe Ratio (shows deduplication rate in <code>bpd jobs</code> command output)
DEDUP_SPACE_RATIO	Dedupe Space
DSTMEDIA_SERVER	Dest Media Svr (writing media server)
DSTMEDIAID	Dest Media ID (writing media ID)
DSTSTORAGE_UNIT	Dest StUnit (writing storage unit)
ELAPSED	Elapsed (elapsed time)
ENDED	Ended
ESTFILE	Est File (estimated number of files)
ESTKB	Est KB (estimated number of kilobytes)
FILES	Files
GROUP	Group
JOBID	JobID
KBPERSEC	KB Per Sec
KILOBYTES	Kilobytes
LASTBACKUP	Last Backup (date and time)
MAINPID	Main PID (PID that spawns job, if applicable)

Table 2-69 COLDEFS entries (*continued*)

COLDEFS entry	Column Name
NUMTAPESEJECT	Media to Eject (number of tapes to eject; Vault only)
OPERATION	Operation (current operation)
OWNER	Owner
PATHNAME	Pathname
PARENTJOBID	Parent JobID
POLICY	Policy
POLICYTYPE	Policy Type
PRIORITY	Priority
PROFILE	Profile (Vault only)
RETENTION	Retention (retention period)
RESUMABLE	Resumable
ROBOT	Robot (Vault only)
RQSTPID	Request PID (PID requesting job, if applicable)
SCHEDULE	Schedule
SCHEDULETYPE	Schedule Type
SESSIONID	Session ID (Vault only)
SRCMEDIA_SERVER	Src Media Svr
SRCMEDIAID	Src Media ID
SRCSTORAGE_UNIT	Src StUnit
STARTED	Started
STATE	State
STATUS	Status
STREAMNUMBER	Stream Number

Table 2-69 COLDEFS entries (continued)

COLDEFS entry	Column Name
SUSPENDABLE	Suspendable
TYPE	Type (job type)
VAULT	Vault (Vault only)

BPDBM_VERBOSE option for NetBackup servers

The `BPDBM_VERBOSE` option is used for debugging purposes. It controls the amount of information NetBackup includes in the `bpdbm` debug log.

Table 2-70 BPDBM_VERBOSE information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	The default is that <code>BPDBM_VERBOSE</code> is the same value as the <code>VERBOSE</code> option (Global logging level). The <code>BPDBM_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPDBM_VERBOSE = -1 0 1 2 3 4 5</pre> This entry should appear only once in the configuration file.

Table 2-70 BPDBM_VERBOSE information (*continued*)

Usage	Description
Example	- To use the same value as the <code>VERBOSE</code> option for, enter: <code>BPDBM_VERBOSE = 0</code> This option is the same as setting the BPDBM logging level to Same as Global in the Logging host properties. - To log the minimum amount of information, enter: <code>BPDBM_VERBOSE = -1</code> This option is the same as setting the BPDBM logging level to 0 in the Logging host properties. - To log additional information, enter a value of 1 through 5: <code>BPDBM_VERBOSE = 1</code> This option is the same as setting the BPDBM logging level to 1 in the Logging host properties. - To log the maximum amount of information, enter: <code>BPDBM_VERBOSE = 5</code> This option is the same as setting the BPDBM logging level to 5 in the Logging host properties. The following examples show two entries which enable logging, while they minimize the growth rate of the <code>bpdbm</code> debug file: <pre>VERBOSE = 5 BPDBM_VERBOSE = -1</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Logging > BPDBM logging level. See the NetBackup Logging Reference Guide for more information about the debug log. See “Logging properties” on page 113.

BPRD_VERBOSE option for NetBackup servers

Used for debugging purposes, the `BPRD_VERBOSE` option controls the amount of information that NetBackup includes in the `bprd` debug logs.

Table 2-71 BPRD_VERBOSE information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-71 BPRD_VERBOSE information (continued)

Usage	Description
How to use	The default is that the value is the same as the <code>VERBOSE</code> option (Global logging level). The <code>BPRD_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPRD_VERBOSE = -1 0 1 2 3 4 5</pre> This entry should appear only once in the configuration file.
Example	■ To use the same value as the <code>VERBOSE</code> option, enter: <pre>BPRD_VERBOSE = 0</pre> This option is the same as setting the BPRD logging level in the Logging host properties to Same as Global. ■ To log the minimum amount of information, enter: <pre>BPRD_VERBOSE = -1</pre> This option is the same as setting the BPRD logging level to 0 in the Logging host properties. ■ To log additional information, enter a value of 1 through 5: <pre>BPRD_VERBOSE = 1</pre> This option is the same as setting the BPRD logging level to 1 in the Logging host properties. ■ To log the maximum amount of information, enter: <pre>BPRD_VERBOSE = 5</pre> This option is the same as setting the BPRD logging level to 5 in the Logging host properties.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Logging > BPRD logging level. See the NetBackup Logging Reference Guide for more information about the debug log. See “Logging properties” on page 113.

BPTM_VERBOSE option for NetBackup servers

The `BPTM_VERBOSE` option is used for debugging purposes. It controls the amount of information that NetBackup includes in the `bptm` debug logs.

Table 2-72 BPTM_VERBOSE information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	The default is that <code>BPTM_VERBOSE</code> is the same value as the <code>VERBOSE</code> option (Global logging level). The <code>BPTM_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPTM_VERBOSE = -1 0 1 2 3 4 5</pre> This entry should appear only once in the configuration file.
Example	- To use the same value as the <code>VERBOSE</code> option, enter: <code>BPTM_VERBOSE = 0</code> This option is the same as setting the BPTM logging level in the Logging host properties to Same as Global. - To log the minimum amount of information, enter: <code>BPTM_VERBOSE = -1</code> This option is the same as setting the BPTM logging level to 0 in the Logging host properties. - To log additional information, enter a value of 1 through 5: <code>BPTM_VERBOSE = 1</code> This option is the same as setting the BPTM logging level to 1 in the Logging host properties. - To log the maximum amount of information, enter: <code>BPTM_VERBOSE = 5</code> This option is the same as setting the BPTM logging level to 5 in the Logging host properties.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Logging > BPTM logging level. See the NetBackup Logging Reference Guide for more information about the debug log. See “Logging properties” on page 113.

BPEND_TIMEOUT option for NetBackup servers

The `BPEND_TIMEOUT` option specifies the number of seconds to wait for the `bpend_notify` script on a client to complete.

Table 2-73 BPEND_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPEND_TIMEOUT = seconds</pre> The default timeout is 300 seconds (five minutes). Note: If this option is changed, verify that the <code>CLIENT_READ_TIMEOUT</code> option is set to the same value or higher. This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Timeouts > Backup end notify timeout. See "Timeouts properties" on page 175.

BPSTART_TIMEOUT option for NetBackup servers

The `BPSTART_TIMEOUT` option specifies the number of seconds to wait for the `bpstart_notify` script on a client to complete.

Table 2-74 BPSTART_TIMEOUT information

Usage	Description
Where to use	On NetBackup media servers.

Table 2-74 BPSTART_TIMEOUT information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPSTART_TIMEOUT = seconds</pre> The default timeout is 300 seconds (five minutes). Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can timeout while the <code>bpstart_notify</code> script is running. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the media server > Timeouts > Backup start notify timeout. See "Timeouts properties" on page 175.

CHECK_RESTORE_CLIENT option for NetBackup servers

The `CHECK_RESTORE_CLIENT` option specifies that the client to be restored to is checked before the restore starts. An unresponsive client can slow restores for other clients that have data on the same tapes.

Table 2-75 CHECK_RESTORE_CLIENT information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. This option should appear only once in the configuration file.
Equivalent host property	No equivalent exists in the host properties.

CLIENT_CONNECT_TIMEOUT option for NetBackup servers

This option specifies the number of seconds that the server waits when it connects to a client. If the server needs to wait longer than the time specified, it times out.

Table 2-76 CLIENT_CONNECT_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CLIENT_CONNECT_TIMEOUT = seconds</pre> The default timeout is 300 seconds (five minutes). This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Timeouts > Client connect timeout. See "Timeouts properties" on page 175.

CLIENT_PORT_WINDOW option for NetBackup servers and clients

This option specifies the range of non-reserved ports on this computer that are used as source ports when connecting to NetBackup on other computers. This setting applies to daemon or service socket connections to the server and to the client hosts as well as call-back from `bpcd`.

Table 2-77 CLIENT_PORT_WINDOW information

Usage	Description
Where to use	On NetBackup servers and clients.

Table 2-77 CLIENT_PORT_WINDOW information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CLIENT_PORT_WINDOW = start_port_range end_port_range</pre> If 0 is specified for the first number (default), the operating system determines the non-reserved port to use. This option should appear only once in the configuration file.
Example	The following example permits ports from 4800 through 5000: <pre>CLIENT_PORT_WINDOW = 4800 5000</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server or client > Port ranges > Client port window. See “Port ranges properties” on page 127.

CLIENT_READ_TIMEOUT option for NetBackup servers

The `CLIENT_READ_TIMEOUT` option specifies the number of seconds that NetBackup waits for a response from a client before the operation attempt fails. For example, if the primary server does not get a response from a client within the `CLIENT_READ_TIMEOUT` period, the backup or restore operation fails.

Table 2-78 CLIENT_READ_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary and media servers.

Table 2-78 CLIENT_READ_TIMEOUT information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CLIENT_READ_TIMEOUT = seconds</pre> By default, <code>CLIENT_READ_TIMEOUT</code> is not present on the server or the database agent and the client-read timeout is 300 seconds (five minutes). This time is a reasonable default. Change only in the event of problems. <code>CLIENT_READ_TIMEOUT</code> on a database agent is a special case because these types of clients can initially require more time to get ready than other clients. Database backup utilities frequently start several backup jobs at the same time, which can slow the CPU. The sequence on a database agent is as follows: ■ NetBackup on the database agent reads the client's <code>CLIENT_READ_TIMEOUT</code> to find the value to use initially. If the option is not set, the standard default of five minutes is used. ■ When the database agent API receives the server's value, it uses it as the <code>CLIENT_READ_TIMEOUT</code>. It is recommended to increase the timeout value on the primary and the media server in the following situations: ■ For database agents, a setting of 15 minutes is adequate for many installations. ■ Backing up directly to an MSDP cloud storage server. If the value is not increased, you may see jobs failing with the following message in the job details: <pre>Error bpbrm (pid=119850) socket read failed: errno = 62 - Timer expired</pre> Note that increasing the timeout is not needed if you use a storage lifecycle policy to first back up to an MSDP storage server and then duplicate the data to an MSDP cloud storage server or an MSDP cloud LSU using an optimized duplication operation. (This operation is the recommended method of operation.) Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can timeout while the <code>bpstart_notify</code> script is running.
Example	The following example configures a client read timeout of 15 minutes. <pre>CLIENT_READ_TIMEOUT = 900</pre>

Table 2-78 CLIENT_READ_TIMEOUT information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or media server > Timeouts > Client read timeout. See “Timeouts properties” on page 175.

CLOUD_AUTODISCOVERY_INTERVAL for NetBackup servers

This option controls how often NetBackup scans the Snapshot Manager servers to discover cloud assets to display in NetBackup.

Table 2-79 CLOUD_AUTODISCOVERY_INTERVAL information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. Note: These commands require administrator privilege on the NetBackup primary server. For assistance, contact the NetBackup administrator. The default is 2 hours. The minimum is 2 hours, the maximum 1 year. Use the following format: <code>CLOUD_AUTODISCOVERY_INTERVAL = number of seconds</code> For example: <code>CLOUD_AUTODISCOVERY_INTERVAL = 100000</code> This entry should appear only once in the configuration file. Note: After changing this option, stop and restart the NetBackup services.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

CLUSTER_ECA_CERT_PATH for clustered primary server

The `CLUSTER_ECA_CERT_PATH` option is specific to clustered primary server. It specifies the path to the external CA-signed certificate of the virtual name.

Table 2-80 CLUSTER_ECA_CERT_PATH information

Usage	Description
Where to use	On clustered primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <code>CLUSTER_ECA_CERT_PATH = Path to the certificate of the virtual identity</code>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

CLUSTER_ECA_KEY_PASSPHRASEFILE for clustered primary server

The `CLUSTER_ECA_KEY_PASSPHRASEFILE` option is specific to clustered primary server. It specifies the path to the text file where the passphrase for the virtual name certificate's private key is stored.

`CLUSTER_ECA_KEY_PASSPHRASEFILE` is optional. You should define this option if the virtual name certificate's private key is encrypted.

See “[CLUSTER_ECA_PRIVATE_KEY_PATH for clustered primary server](#)” on page 206.

Table 2-81 CLUSTER_ECA_KEY_PASSPHRASEFILE information

Usage	Description
Where to use	On clustered primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <code>CLUSTER_ECA_KEY_PASSPHRASE_FILE = Path to the passphrase file</code>

Table 2-81 CLUSTER_ECA_KEY_PASSPHRASEFILE information
(continued)

Usage	Description
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

CLUSTER_ECA_PRIVATE_KEY_PATH for clustered primary server

The `CLUSTER_ECA_PRIVATE_KEY_PATH` option is specific to clustered primary server. It specifies the path to the private key for the external CA-signed certificate of the virtual name.

If the virtual name certificate's private key is encrypted, you should define the `CLUSTER_ECA_KEY_PASSPHRASEFILE` option.

See “[CLUSTER_ECA_KEY_PASSPHRASEFILE for clustered primary server](#)” on page 205.

Table 2-82 CLUSTER_ECA_PRIVATE_KEY_PATH information

Usage	Description
Where to use	On clustered primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CLUSTER_ECA_PRIVATE_KEY_PATH = Path to the private key of the external certificate</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

CLUSTER_ECA_TRUST_STORE_PATH for clustered primary server

The `CLUSTER_ECA_TRUST_STORE_PATH` option is specific to clustered primary server. It specifies the path to the certificate bundle file that contains all trusted root CA certificates in PEM format.

Table 2-83 CLUSTER_ECA_TRUST_STORE_PATH information

Usage	Description
Where to use	On clustered primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CLUSTER_ECA_TRUST_STORE_PATH = Path to the external CA certificate</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

COMPUTE_IMAGE_ENTROPY for NetBackup primary servers

Use the `COMPUTE_IMAGE_ENTROPY` option to enable computation of entropy and file attributes in NetBackup that enhances cyber resiliency in a NetBackup - Helios environment.

The entropy metric is used with the anomaly detection in Helios to help you detect potential malicious activity.

Table 2-84 COMPUTE_IMAGE_ENTROPY information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-84 COMPUTE_IMAGE_ENTROPY information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>COMPUTE_IMAGE_ENTROPY = Value</pre> You can specify one of the following values for the <code>COMPUTE_IMAGE_ENTROPY</code> option: ■ <code>ALWAYS</code> - Computation of entropy and file attributes is always enabled. This is the default value. ■ <code>NEVER</code> - Computation of entropy and file attributes is always disabled. ■ <code>IF_MANAGED_BY_ALTA</code> - Computation of entropy and file attributes is enabled if Helios manages the associated NetBackup primary server. If Helios does not manage the primary server, computation is disabled. Note: After the NetBackup primary server is registered with the Helios server, computation of entropy and file attributes starts within the next 24 hours with the new backup jobs.
Equivalent NetBackup web UI property	No equivalent exists.

CONNECT_OPTIONS option for NetBackup servers

The `CONNECT_OPTIONS` apply to connections to the local host only, as follows (they do *not* apply to connections to remote hosts):

- Whether subsequent call-back connections with *host* use the traditional call-back method, use `vnetd`, or use a `PBX/vnetd` forwarding connection.
- Whether connections to *host* use reserved or a non-reserved source port number.

Table 2-85 CONNECT_OPTIONS information

Usage	Description
Where to use	On NetBackup primary servers or media servers.

Table 2-85 CONNECT_OPTIONS information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CONNECT_OPTIONS = host [0 1 2] [0 1 2]</pre> The following variables are defined: <i>Host</i> is a host name local to this host. You may have multiple <code>CONNECT_OPTIONS</code> entries in the configuration, and <code>localhost</code> overrides other local host names. If a local host name is not specified in any <code>CONNECT_OPTIONS</code> entries, the values from the <code>DEFAULT_CONNECT_OPTIONS</code> option are used. See "DEFAULT_CONNECT_OPTIONS option for NetBackup servers" on page 211.
	The first setting indicates the type of port to use as the source port for connections to service daemons on <i>host</i>: 0 = Connections on this computer should be from a reserved source port number. 1 = Connections on this computer should be from a non-reserved source port number that is selected from the <code>CLIENT_PORT_WINDOW</code> range. (The default is 1.) In the NetBackup web UI, open the media server host properties and select Universal settings. 2 = Use the method that the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option defines. See "Universal settings properties" on page 178.
	The second setting indicates the call-back method to use with <i>host</i>. (This method applies if <code>bpcd</code> cannot be reached using ports 1556 or 13724.) 0 = Use the traditional call-back method. <i>Host</i> connects back to a random port number that this computer has selected from the <code>SERVER_RESERVED_PORT_WINDOW</code> range, or <code>SERVER_PORT_WINDOW</code> range as determined by the first setting. 1 = Use the <code>vnetd</code> no call-back method. Connect to <code>vnetd</code> instead of a random port. Attempt to connect to port 1556 before attempting to connect to <code>vnetd</code>. 2 = Use the method that the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option defines (default).

Table 2-85 CONNECT_OPTIONS information (*continued*)

Usage	Description
Example 1	The configuration file can contain <code>CONNECT_OPTIONS</code> settings for local host names. <pre>CONNECT_OPTIONS = localhost 0 0</pre> In this example, local connections to daemons on the local host <code>shark</code> attempt to use port 1556. If the previous attempt was unsuccessful, then the connections try <code>vnetd</code>. If the connections are successful using 1556 or <code>vnetd</code>, then both settings are ignored. <pre>\$ bptestbpcd -host shark 0 0 10.82.105.11:40402 -> 10.82.105.11:1556 10.82.105.11:40404 -> 10.82.105.11:1556</pre>
Example 2	<pre>CONNECT_OPTIONS = host 0 1</pre> In this example: ■ Call-back connections are to <code>vnetd</code> on this computer. ■ The source ports for the daemon connection are bound from the reserved port number range.
Example 3	<pre>CONNECT_OPTIONS = host 1 1</pre> In this example: ■ Call-back connections are to <code>vnetd</code> on this computer. ■ The source ports for the daemon connection are bound from the non-reserved port number range.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or media server > Firewall. See "Firewall properties" on page 104.

DATAACCESS_AUDIT_INTERVAL_HOURS for NetBackup primary servers

Use the `DATAACCESS_AUDIT_INTERVAL_HOURS` option to set an interval to periodically add audit records for the browse image (`bplist`) operations into the NetBackup database.

Consider the following example:

The `DATAACCESS_AUDIT_INTERVAL_HOURS` option is set to 2 hours. All the audit records for the `bplist` operations are cached for 2 hours. One of the many similar `bplist` audit records is identified and is added into the database every 2 hours.

This option prevents the database size from increasing exponentially because of the `bplist` audit records.

To add all the `bplist` audit records from the cache into the NetBackup database, run the following command on the primary server:

```
nbcertcmd -postAudit -dataAccess
```

Table 2-86 DATAACCESS_AUDIT_INTERVAL_HOURS information

Usage	Description
Where to use	On primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GENERIC_KEY_VAL_LIST = (DATAACCESS_AUDIT_INTERVAL_HOURS) (time in hours)</pre> The default interval is 1 hour.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

DEFAULT_CONNECT_OPTIONS option for NetBackup servers

The `DEFAULT_CONNECT_OPTIONS` option specifies the default values for the `CONNECT_OPTIONS` configuration option. If a host name is not specified in any `CONNECT_OPTIONS` option, the value from the `DEFAULT_CONNECT_OPTIONS` option is used.

Note: The `DEFAULT_CONNECT_OPTIONS` apply to connections to the local host only; they do *not* apply to connections to remote hosts.

See “[CONNECT_OPTIONS option for NetBackup servers](#)” on page 208.

Table 2-87 DEFAULT_CONNECT_OPTIONS information

Usage	Description
Where to use	On NetBackup primary servers or media servers.

Table 2-87 DEFAULT_CONNECT_OPTIONS information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DEFAULT_CONNECT_OPTIONS = [0 1][0 1]</pre> The default value is 0 1. This option should appear only once in the configuration file.
	The first setting indicates the type of port to use as the source port when connecting to the <code>bpcd</code> daemon port on the local host. It also indicates the type of server port if using the traditional call-back method. 0 = Connections on this computer should use a reserved port number. They are selected from the <code>SERVER_RESERVED_PORT_WINDOW</code> range if using the traditional call-back method. 1 = Connections on this computer should use a non-reserved port number. Connections are selected from the <code>CLIENT_PORT_WINDOW</code> range for source ports and from the <code>SERVER_PORT_WINDOW</code> range if using the traditional call-back method.
	The second setting indicates the call-back method to use. (This setting applies if <code>bpcd</code> cannot be reached using ports 1556 or 13724.) 0 = Use the traditional call-back method. The destination host connects back to a random port number that this computer has selected from the <code>SERVER_RESERVED_PORT_WINDOW</code> range, or the <code>SERVER_PORT_WINDOW</code> range as determined by the first setting. 1 = Use the <code>vnetd</code> no call-back method. Connect to <code>vnetd</code> instead of a random port. Attempt to connect to port 1556 before attempting to connect to <code>vnetd</code>.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or media server > Firewall. See "Firewall properties" on page 104.

DISABLE_CERT_AUTO_RENEW option for NetBackup servers and clients

This option disables the automatic renewal of host ID-based certificates.

For more information about the automatic renewal of host ID-based certificates, see the [NetBackup Security and Encryption Guide](#).

Table 2-88 DISABLE_CERT_AUTO_RENEW information

Usage	Description
Where to use	On NetBackup primary servers and clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DISABLE_CERT_AUTO_RENEW = 1</pre> This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

DISABLE_JOB_LOGGING option for NetBackup servers

This option disables the logging of the job information that the NetBackup Activity Monitor requires.

Table 2-89 DISABLE_JOB_LOGGING information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DISABLE_JOB_LOGGING</pre> The default is that this option is not present in the configuration file and that job logging occurs.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Media > Enable job logging. See "Media properties" on page 119.

DISABLE_STANDALONE_DRIVE_EXTENSIONS option for NetBackup servers

This option disables the nonrobotic drive operations. During a backup, NetBackup automatically attempts to use standalone volumes in nonrobotic drives.

Table 2-90 DISABLE_STANDALONE_DRIVE_EXTENSIONS information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbemmcmd</code> command to change the option. For example: <pre>nbemmcmd -changesetting -DISABLE_STANDALONE_DRIVE_EXTENSIONS no</pre> For information about these commands, see the NetBackup Commands Reference Guide. This option should appear only once in the configuration file.
Example	The following command enables nonrobotic drive operations. <pre>nbemmcmd -changesetting -DISABLE_STANDALONE_DRIVE_EXTENSIONS no</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Media > Enable standalone drive extension. The default is that this option is enabled. See "Media properties" on page 119.

DISALLOW_BACKUPS_SPANNING_MEDIA option for NetBackup servers

This option prevents backups from spanning media.

Table 2-91 DISALLOW_BACKUPS_SPANNING_MEDIA information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-91 DISALLOW_BACKUPS_SPANNING_MEDIA information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DISALLOW_BACKUPS_SPANNING_MEDIA</pre> The default is that the entry is not present in the configuration file and backups are allowed to span media.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Media > Allow backups to span tape media. See "Media properties" on page 119.

DISALLOW_CLIENT_LIST_RESTORE option for NetBackup servers

This option denies the list and restore requests for all clients. When this option is present, clients cannot list or restore any files that they have backed up through this primary server.

Table 2-92 DISALLOW_CLIENT_LIST_RESTORE information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DISALLOW_CLIENT_LIST_RESTORE</pre> The default is that the entry is not present in the configuration file and clients can list and restore their files. Note: Override the <code>DISALLOW_CLIENT_LIST_RESTORE</code> option for individual clients by changing their <code>list_restore</code> setting.

Table 2-92 DISALLOW_CLIENT_LIST_RESTORE information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Client attributes > Allow client restore. See “Client attributes properties” on page 64.

DISALLOW_CLIENT_RESTORE option for NetBackup servers

This option denies the restore requests for all clients. When this option is present, clients cannot restore the files that they have backed up through this primary server.

Table 2-93 DISALLOW_CLIENT_RESTORE information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DISALLOW_CLIENT_RESTORE</pre> The default is that the entry is not present in the configuration file and clients can restore their files. Note: To override the <code>DISALLOW_CLIENT_RESTORE</code> option for individual clients, change their <code>list_restore</code> setting.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Client attributes > Allow client browse. See “Client attributes properties” on page 64.

DISALLOW_SERVER_FILE_WRITES option for NetBackup servers and clients

The `DISALLOW_SERVER_FILE_WRITES` entry prevents the NetBackup server from creating files on the NetBackup server or client. This entry prevents NetBackup servers from remotely performing restores or remotely changing client configurations.

For further information on the commands that are described in the following tables, see the [NetBackup Commands Reference Guide](#).

Table 2-94 DISALLOW_SERVER_FILE_WRITES information

Where to use	How to use	Notes when used on a local host	Notes when used remotely	Default behavior
NetBackup web UI	Hosts > Host properties > Select host > Universal settings > Allow server file writes	Allow server file writes can be set on an all-in-one host that contains the primary server, media server, and client. If the media server and client are not on the primary server, you must use the DISALLOW_SERVER_FILE_WRITES entry on the media server or client.	Allow server file writes cannot be set from the NetBackup web UI . On the media server or client, use the DISALLOW_SERVER_FILE_WRITES entry. See Table 2-96 .	The server writes are allowed.
NetBackup Backup, Archive, and Restore Windows client interface	File > NetBackup Client Properties > Allow server-directed restores	Allow server-directed restores can be used only from a Windows computer.	Allow server-directed restores cannot be used remotely. On the media server or client, use the DISALLOW_SERVER_FILE_WRITES entry. See Table 2-96 .	The server-directed restores are allowed.
bpsetconfig or bpgetconfig (use only on primary server or media server) nbsetconfig or nbgetconfig (use on primary server, media server, or client)	For command examples, see Table 2-95 and Table 2-96 .	DISALLOW_SERVER_FILE_WRITES can be enabled and disabled.	bpsetconfig and bpgetconfig can be run only from the primary server or media server. Note: DISALLOW_SERVER_FILE_WRITES = No using bpsetconfig or nbsetconfig cannot be set remotely (can only be set locally).	DISALLOW_SERVER_FILE_WRITES = No

Table 2-95 `bpsetconfig` and `nbsetconfig` examples for a local host

Command	Command examples for a local host
<code>bpsetconfig</code>	From the local primary server or media server On Windows: <pre>install_path\NetBackup\bin\admincmd>bpsetconfig</pre> <pre>bpsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>bpsetconfig> <ctl-Z></pre> On UNIX: <pre>/usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre> <pre>bpsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>bpsetconfig> <ctl-D></pre>
<code>nbsetconfig</code>	From the local primary server, media server, or client On Windows: <pre>install_path\NetBackup\bin>nbsetconfig</pre> <pre>nbsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>nbsetconfig> <ctl-Z></pre> On UNIX: <pre>/usr/opensv/netbackup/bin/nbsetconfig</pre> <pre>nbsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>nbsetconfig> <ctl-D></pre>

Table 2-96 `bpsetconfig` and `nbsetconfig` examples for a remote host

Command	Command examples for a remote host
<code>bpsetconfig</code>	From the remote primary server or media server On Windows: <pre>install_path\NetBackup\bin\admincmd>bpsetconfig -h host</pre> <pre>bpsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>bpsetconfig> <ctl-Z></pre> On UNIX: <pre>/usr/opensv/netbackup/bin/admincmd/bpsetconfig -h host</pre> <pre>bpsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>bpsetconfig> <ctl-D></pre>
<code>nbsetconfig</code>	From the remote primary server, media server, or client On Windows: <pre>install_path\NetBackup\bin>nbsetconfig -h host</pre> <pre>nbsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>nbsetconfig> <ctl-Z></pre> On UNIX: <pre>/usr/opensv/netbackup/bin/nbsetconfig -h host</pre> <pre>nbsetconfig> DISALLOW_SERVER_FILE_WRITES = Yes</pre> <pre>nbsetconfig> <ctl-D></pre>

DTE_IGNORE_IMAGE_MODE for NetBackup servers

Use the `DTE_IGNORE_IMAGE_MODE` option if you do not want the data to be encrypted even if the data-in-transit encryption (DTE) mode of the backup image is enabled.

The `DTE_IGNORE_IMAGE_MODE` option is applicable for all backup images.

Table 2-97 `DTE_IGNORE_IMAGE_MODE` information

Usage	Description
Where to use	On NetBackup servers.

Table 2-97 DTE_IGNORE_IMAGE_MODE information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DTE_IGNORE_IMAGE_MODE = NEVER ALWAYS WHERE_UNSUPPORTED</pre> The default value of the <code>DTE_IGNORE_IMAGE_MODE</code> option is <code>NEVER</code>. ■ <code>NEVER</code> - Use this option to specify that the data-in-transit encryption takes place based on the DTE mode of the image. ■ <code>ALWAYS</code> - Use this option to specify that the DTE mode of the image is always ignored during data-in-transit encryption irrespective of whether the NetBackup host supports the encryption or not. Data-in-transit encryption takes place based on the global DTE mode and client DTE mode. ■ <code>WHERE_UNSUPPORTED</code> - Use this option if you have NetBackup hosts earlier than 9.1 in your environment and you do not want the jobs to fail for these hosts when the DTE mode is enabled for the image. With this configuration, data-in-transit encryption happens based on the global and client DTE mode settings. The image DTE mode is ignored.
Equivalent NetBackup web UI property	No equivalent exists.

ECA_CERT_PATH for NetBackup servers and clients

The `ECA_CERT_PATH` option specifies the path to the external CA-signed certificate of the host. This option is mandatory.

NetBackup supports the following certificate sources for host certificates:

- Windows certificate store

Note: The Windows certificate store is not supported for clustered primary servers.

- File-based certificates

Certificate order in the certificate file

A certificate file must have a certificate chain with certificates in the correct order. The chain starts with the server certificate (also known as the leaf certificate) and is followed by zero or more intermediate certificates. The chain must contain all intermediate certificates up to the Root CA certificate but should not contain the Root CA certificate itself. The chain is created such that each certificate in the chain signs the previous certificate in the chain.

The certificate file should be in one of the following formats:

- PKCS #7 or P7B file that is either DER or PEM encoded that has certificates in the specified order
- A file with the PEM certificates that are concatenated together in the specified order

Table 2-98 ECA_CERT_PATH information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. For file-based certificates, use the following format: <i>ECA_CERT_PATH = Path to the external certificate of the host</i> For example: <code>c:\server.pem</code> If you use this option on a Flex Appliance application instance, the path must be <code>/mnt/nbdata/hostcert/</code>. For Windows certificate store, use the following format: <i>ECA_CERT_PATH = Certificate store name\Issuer name\Subject name</i> You can specify multiple certificate selection queries in a comma-separated format. <i>ECA_CERT_PATH = Store name1\Issuer name1\Subject name1,Store name2\Issuer name2\Subject name2</i> See "Specifying Windows certificate store for ECA_CERT_PATH" on page 222.

Table 2-98 ECA_CERT_PATH information (continued)

Usage	Description
Equivalent NetBackup web UI property	No equivalent exists.

Specifying Windows certificate store for ECA_CERT_PATH

NetBackup selects a certificate from any of the local machine certificate stores on a Windows host.

In case of Windows certificate store, ECA_CERT_PATH is a list of comma-separated clauses.

Each clause is of the form *Store name\Issue\Subject*. Each clause element contains a query.

`$hostname` is a keyword that is replaced with the fully qualified domain name of the host. Use double quotes when a `\` is present in the actual path. For example, `MY\Veritas\"NetBackup\"$hostname"`.

`$shorthostname` is a keyword that is replaced with the short name of the host. Use double quotes when a `\` is present in the actual path. For example, `MY\Veritas\"NetBackup\"$shorthostname"`.

The 'Store name' should be the exact name of the store where the certificate resides. For example: 'MY'

The 'Issuer' is optional. If this is provided, NetBackup picks the certificates for which the Issuer DN contains the provided substring.

The 'Subject' is mandatory. NetBackup picks the certificate for which the Subject DN contains the provided substring.

You must ensure to:

- Add the root certificate to Trusted Root Certification Authorities or Third-Party Root Certification Authorities in the Windows certificate store.
- If you have any intermediate CAs, add their certificates to the Intermediate Certification Authorities in the Windows certificate store.

Example - Certificate locations with WHERE CLAUSE:

- `My\Veritas\"$hostname, My\ExampleCompany\"$hostname`
Where (certificate store is MY, Issuer DN contains `Veritas`, Subject DN contains `$hostname`) OR (certificate store name is MY, Issuer DN contains `ExampleCompany`, Subject DN contains `$hostname`)

- `MY\Veritas\NetBackup\%hostname"`
 Where certificate store name is `MY`, Issuer DN contains `Veritas`, Subject DN contains `NetBackup\%hostname`
- `MY\\%hostname`
 Where certificate store name is `MY`, any Issuer DN, Subject DN contains `%hostname`
- `MY\\%shorthostname`
 Where certificate store name is `MY`, any Issuer DN, Subject DN contains `%shorthostname`
- `MY\Veritas\NetBackup %hostname`
 Where certificate store name is `MY`, Issuer DN contains `Veritas`, Subject DN contains `NetBackup %hostname`

If you provide a space between words, it is considered as a valid character.

Example - Certificate locations with invalid data:

- `MY\`
 The Subject DN should have some value.
- `My\%hostname`
 The Subject DN should have some value.
- `\\%hostname`
 The certificate store name should have exact value of the store in which the certificate resides.
- `MY\CN=Veritas\CN=%hostname`
 The Subject DN and issuer DN cannot contain `=`, and also specific tags like `CN=`.

ECA_CRL_CHECK for NetBackup servers and clients

The `ECA_CRL_CHECK` option lets you specify the revocation check level for external certificates of the host. It also lets you disable the revocation check for the external certificates. Based on the check, revocation status of the certificate is validated against the Certificate Revocation List (CRL) during host communication.

You can choose to use the CRLs from the directory that is specified for the `ECA_CRL_PATH` configuration option in the configuration file (`bp.conf` on UNIX or Windows registry) or the CRL Distribution Point (CDP).

See [“ECA_CRL_PATH for NetBackup servers and clients”](#) on page 224.

Table 2-99 ECA_CRL_CHECK information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ECA_CRL_CHECK = CRL check</pre> You can specify one of the following: ■ DISABLE (or 0) - Revocation check is disabled. Revocation status of the certificate is not validated against the CRL during host communication. ■ LEAF (or 1) - Revocation status of the leaf certificate is validated against the CRL. This is the default value. ■ CHAIN (or 2) - Revocation status of all certificates from the certificate chain are validated against the CRL.
Equivalent web UI property	No equivalent exists.

ECA_CRL_PATH for NetBackup servers and clients

The `ECA_CRL_PATH` option specifies the path to the directory where the Certificate Revocation Lists (CRL) of the external certificate authority (ECA) are located.

These CRLs are copied to NetBackup CRL cache. Revocation status of the external certificate is validated against the CRLs from the CRL cache.

CRL in the CRL cache is periodically updated with the CRL on the location that is specified for `ECA_CRL_PATH` based on the `ECA_CRL_PATH_SYNC_HOURS` option.

If the `ECA_CRL_CHECK` or `HADOOP_CRL_CHECK` option is not set to `DISABLE` (or 0) and the `ECA_CRL_PATH` option is not specified, NetBackup downloads the CRLs from the URLs that are specified in the CRL distribution point (CDP) and uses them to verify revocation status of the peer host's certificate.

Note: For validating the revocation status of a virtualization server certificate, the `VIRTUALIZATION_CRL_CHECK` option is used.

See “[VIRTUALIZATION_CRL_CHECK for NetBackup servers and clients](#)” on page 322.

For validating the revocation status of a Hadoop server certificate, the `HADOOP_CRL_CHECK` option is used.

Table 2-100 `ECA_CRL_PATH` information

Usage	Description
Where to use	On NetBackup servers or clients. If certificate validation is required for VMware, Red Hat Virtualization servers, Nutanix AHV, or Hadoop, this option must be set on the NetBackup primary server and respective access or backup hosts, irrespective of the certificate authority that NetBackup uses for host communication (NetBackup CA or external CA). If certificate validation is required for VMware, Red Hat Virtualization servers, or Hadoop, this option must be set on the NetBackup primary server and respective access or backup hosts, irrespective of the certificate authority that NetBackup uses for host communication (NetBackup CA or external CA).
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format to specify a path to the CRL directory: <pre>ECA_CRL_PATH = Path to the CRL directory</pre> For example: <pre>ECA_CRL_PATH = /usr/eca/crl/eca_crl_file.crl</pre> If you use this option on a Flex Appliance application instance, the path must be <code>/mnt/nbdata/hostcert/crl</code>.
Equivalent UI property	No equivalent exists.

ECA_CRL_PATH_SYNC_HOURS for NetBackup servers and clients

The `ECA_CRL_PATH_SYNC_HOURS` option specifies the time interval in hours to update the Certificate revocation lists (CRL) in the NetBackup CRL cache with the CRLs in the directory that is specified for the `ECA_CRL_PATH` configuration option.

See [“ECA_CRL_PATH for NetBackup servers and clients”](#) on page 224.

The `ECA_CRL_PATH_SYNC_HOURS` option is not applicable if CDP is used for CRLs.

By default, CRLs in the cache are updated every one hour.

During host communication, revocation status of the external certificate is validated against the CRLs from the CRL cache.

Table 2-101 `ECA_CRL_PATH_SYNC_HOURS` information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ECA_CRL_PATH_SYNC_HOURS = Number of hours</pre> Minimum number of hours that you can specify - 1 hour Maximum number of hours that you can specify - 720 hour The default value is one hour.
Equivalent UI property	No equivalent exists.

ECA_CRL_REFRESH_HOURS for NetBackup servers and clients

The `ECA_CRL_REFRESH_HOURS` option specifies the time interval in hours to download the CRLs from the URLs that are specified in the peer host certificate's CRL distribution points (CDP).

The `ECA_CRL_REFRESH_HOURS` option is applicable when you use CDP for CRLs.

See [“ECA_CRL_PATH for NetBackup servers and clients”](#) on page 224.

After the specified time interval, CRLs of the certificate authority are downloaded from the URLs that are available in CDP.

By default, the CRLs are downloaded from the CDP after every 24 hours.

Table 2-102 ECA_CRL_REFRESH_HOURS information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ECA_CRL_REFRESH_HOURS = Number of hours</pre> Minimum number of hours that you can specify - 0 hour, which indicates that CRLs from the CDP are not periodically downloaded. Maximum number of hours that you can specify - 4380 hours The default value for the option is 24 hours. Note: CRLs are also downloaded from the CDP during host communication if they are expired or not available in the CRL cache, irrespective of the time interval set for the <code>ECA_CRL_REFRESH_HOURS</code> option.
Equivalent UI property	No equivalent exists.

ECA_DISABLE_AUTO_ENROLLMENT for NetBackup servers and clients

When NetBackup is configured to use the certificates that an external CA has signed, such certificates are automatically enrolled with the primary server during host communication. If you want to disable automatic enrollment of such certificates, set the `ECA_DISABLE_AUTO_ENROLLMENT` to '1'.

When automatic enrollment is disabled, you can enroll the external certificates manually using the `nbcertcmd -enrollCertificate` command.

A certificate must be enrolled with the primary server before it can be used for host communication.

By default, automatic certificate enrollment is enabled.

Table 2-103 ECA_DISABLE_AUTO_ENROLLMENT information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ECA_DISABLE_AUTO_ENROLLMENT = 1</pre>
Equivalent UI property	No equivalent exists.

ECA_DR_BKUP_WIN_CERT_STORE for NetBackup servers and clients

The `ECA_DR_BKUP_WIN_CERT_STORE` option specifies whether you want to take a backup of the Windows certificate store information during catalog backup or not. By default, Windows certificate store information is backed up during catalog backup.

Note: If the Windows certificate store information is not exportable, it cannot be backed up during catalog backup.

Table 2-104 ECA_DR_BKUP_WIN_CERT_STORE information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. If you do not want the catalog backup operation to take a backup of the Windows certificate store information, use the following format: <pre>ECA_DR_BKUP_WIN_CERT_STORE = NO</pre>
Equivalent UI property	No equivalent exists.

ECA_KEY_PASSPHRASEFILE for NetBackup servers and clients

The `ECA_KEY_PASSPHRASEFILE` option specifies the path to the text file where the passphrase for the external certificate's private key is stored.

You should specify the `ECA_KEY_PASSPHRASEFILE` option only if the certificate's private key is encrypted.

See [“ECA_PRIVATE_KEY_PATH for NetBackup servers and clients”](#) on page 229.

Note: You should not specify the `ECA_KEY_PASSPHRASEFILE` option if you use Windows certificate store.

See [“ECA_CERT_PATH for NetBackup servers and clients”](#) on page 220.

Note: Do not use the `ECA_KEY_PASSPHRASEFILE` on the MSDP servers that are used for MSDP direct cloud tiering as it is not supported with MSDP direct cloud tiering.

Table 2-105 `ECA_KEY_PASSPHRASEFILE` information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ECA_KEY_PASSPHRASEFILE = Path to the passphrase file</pre>
Equivalent UI property	No equivalent exists.

ECA_PRIVATE_KEY_PATH for NetBackup servers and clients

The `ECA_PRIVATE_KEY_PATH` option specifies the file path to the private key for the external CA-signed certificate of the host.

This option is mandatory for file-based certificates.

If the private key of the certificate is encrypted, you should specify the `ECA_KEY_PASSPHRASEFILE` option.

See [“ECA_KEY_PASSPHRASEFILE for NetBackup servers and clients”](#) on page 229.

NetBackup supports PKCS #1 and PKCS #8 formatted private keys that are either plain text or encrypted. These may either be PEM or DER encoded. However, if it is PKCS #1 encrypted, it must be PEM encoded.

For encrypted private keys, NetBackup supports the following encryption algorithms:

- DES, 3DES, and AES if the private key is in the PKCS #1 format
- DES, 3DES, AES, RC2, and RC4 if the private key is in the PKCS #8 format

Note: You should not specify the `ECA_PRIVATE_KEY_PATH` option if Windows certificate store is specified for the `ECA_CERT_PATH` option.

See “[ECA_CERT_PATH for NetBackup servers and clients](#)” on page 220.

Table 2-106 `ECA_PRIVATE_KEY_PATH` information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <code>ECA_PRIVATE_KEY_PATH = Path to the private key of the external certificate</code> For example: <code>c:\key.pem</code> If you use this option on a Flex Appliance application instance, the path must be <code>/mnt/nbdata/hostcert/</code>.
Equivalent UI property	No equivalent exists.

ECA_TRUST_STORE_PATH for NetBackup servers and clients

The `ECA_TRUST_STORE_PATH` option specifies the file path to the certificate bundle file that contains all trusted root CA certificates.

This certificate file should have one or more certificates in PEM format.

Do not specify the `ECA_TRUST_STORE_PATH` option if you use the Windows certificate store.

The trust store supports certificates in the following formats:

- PKCS #7 or P7B file having certificates of the trusted root certificate authorities that are bundled together. This file may either be PEM or DER encoded.
- A file containing the PEM encoded certificates of the trusted root certificate authorities that are concatenated together.

This option is mandatory for file-based certificates.

The root CA certificate in Cloudera distribution can be obtained from the Cloudera administrator. It may have a manual TLS configuration or an Auto-TLS enabled for the Hadoop cluster. For both cases, NetBackup needs a root CA certificate from the administrator.

The root CA certificate from the Hadoop cluster can validate the certificates for all nodes and allow NetBackup to run the backup and restore process in case of the secure (SSL) cluster. This root CA certificate is a bundle of certificates that has been issued to all such nodes.

Certificate from root CA must be configured under `ECA_TRUST_STORE_PATH` in case of self-signed, third party CA or Local/Intermediate CA environments. For example: In case of AUTO-TLS enabled Cloudera environments, you can typically find the root CA file named with `cm-auto-global_cacerts.pem` at path `/var/lib/cloudera-scm-agent/agent-cert`. For more details, refer Cloudera documentation.

Table 2-107 `ECA_TRUST_STORE_PATH` information

Usage	Description
Where to use	On NetBackup servers or clients. If certificate validation is required for VMware, Red Hat Virtualization servers, or Nutanix AHV, this option must be set on the NetBackup primary server and respective access hosts, irrespective of the certificate authority that NetBackup uses for host communication (NetBackup CA or external CA).

Table 2-107 ECA_TRUST_STORE_PATH information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ECA_TRUST_STORE_PATH = Path to the external CA certificate</pre> For example: <code>c:\rootCA.pem</code> If you use this option on a Flex Appliance application instance, the path must be <code>/mnt/nbdata/hostcert/</code>.
Equivalent UI property	No equivalent exists.

EAT_VERBOSE option for NetBackup servers and clients

Used for debugging purposes, the `EAT_VERBOSE` option controls the amount of information NetBackup includes in the authentication service (AT) client logs that pertain to NetBackup processes.

Table 2-108 EAT_VERBOSE information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 2-108 EAT_VERBOSE information (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>EAT_VERBOSE = [0 1 2 3 4]</pre> The default is that the entry is not present in the configuration file. The AT logs are generated for the NetBackup processes based on the default logging level, which reports only errors. Following are some of the NetBackup processes that use the AT service: ■ <code>bpnbat</code> ■ <code>bpnbaz</code> ■ <code>nbcertcmd</code> ■ <code>nbsl</code> If the <code>EAT_VERBOSE</code> entry is present in the configuration file, the verbosity of AT logs for the NetBackup processes is based on the <code>EAT_VERBOSE</code> option. The AT logs are stored in the respective process log files. To disable AT logging for NetBackup processes, set the <code>EAT_VERBOSE</code> option to -2 in the configuration file. Use the following format: <pre>EAT_VERBOSE = -2</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

ECA_WIN_CERT_STORE_TIME_LAG_MINUTES for NetBackup servers and clients

Use the `ECA_WIN_CERT_STORE_TIME_LAG_MINUTES` option to work around the communication failure issue that may occur when the server system time and client system time do not match. Because of this time difference, the secure connection may not be established between the two communicating hosts as the security certificate may not be valid yet. The current system time may be behind the 'Valid from' time of the selected certificate.

Table 2-109 ECA_WIN_CERT_STORE_TIME_LAG_MINUTES information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GENERIC_KEY_VAL_LIST = (ECA_WIN_CERT_STORE_TIME_LAG_MINUTES) (time in minutes)</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

ECMS_HOSTS_SECURE_CONNECT_ENABLED for servers

The **ECMS_HOSTS_SECURE_CONNECT_ENABLED** option enables or disables the Host Name verification of the external CMS server during an SSL connection. Before you enable the option, review the 'Configure External Credentials' section in the *NetBackup Administrator's Guide, Volume I*.

By default, the **ECMS_HOSTS_SECURE_CONNECT_ENABLED** option is set to **YES** (Enabled). When enabled, the certificate deployed on the external CMS server (For example, CyberArk Server) must have Common Name or Subject Alternative Name that matches the host name of the external CMS server. Else, the SSL connection to the server fails. The host name verification can be disabled by setting the value of **ECMS_HOSTS_SECURE_CONNECT_ENABLED** option to **NO** or **FALSE**.

Note: Hostname verification involves a server identity check to ensure that the client is talking to the correct server and has not been redirected by a man in the middle attack. The check involves viewing the certificate sent by the server, and verifying that the `dnsName` in the `subjectAltName` field of the certificate matches the host portion of the URL used to make the request.

Table 2-110 ECMS_HOSTS_SECURE_CONNECT_ENABLED information

Usage	Description
Where to use	On NetBackup primary server.

Table 2-110 ECMS_HOSTS_SECURE_CONNECT_ENABLED information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format to disable certificate validation for external CMS servers: <pre>ECMS_HOSTS_SECURE_CONNECT_ENABLED = NO</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

ENABLE_CRITICAL_PROCESS_LOGGING for NetBackup servers and clients

The `ENABLE_CRITICAL_PROCESS_LOGGING` option lets you automatically log critical NetBackup processes. Log directories for the critical processes are created and logging begins when this option is enabled in the **Logging** host properties.

Table 2-111 ENABLE_CRITICAL_PROCESS_LOGGING information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view the option. For information about these commands, see the NetBackup Commands Reference Guide. Initially, the <code>bp.conf</code> file does not contain <code>ENABLE_CRITICAL_PROCESS_LOGGING</code> entry. After logging for critical processes is enabled, a corresponding entry is added in the <code>bp.conf</code> file as follows: <pre>ENABLE_CRITICAL_PROCESS_LOGGING = YES</pre> Note: You should not modify the <code>ENABLE_CRITICAL_PROCESS_LOGGING</code> parameter. To disable the logging for critical processes, modify the logging levels for those processes.

Table 2-111 ENABLE_CRITICAL_PROCESS_LOGGING information
(continued)

Usage	Description
Equivalent web UI host property	Hosts > Host properties > Select the hosts > Logging > Logging for critical processes. See “Logging properties” on page 113.

ENABLE_DIRECT_CONNECTION for servers

With NAT clients in place, NetBackup primary servers and media servers are configured only to accept communication requests from clients.

Servers cannot communicate directly with clients. The `ENABLE_DIRECT_CONNECTION` option lets you establish a direct connection between servers and clients when required.

Here are some example scenarios where servers need to directly connect to clients:

- When the NetBackup domain consists of clients that are not behind any firewall or are not using any gateway
- When the NetBackup domain consists of clients with earlier versions

By default the direct connection between servers and clients is disabled.

Table 2-112 ENABLE_DIRECT_CONNECTION information

Usage	Description
Where to use	On NetBackup servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide . To enable a direct connection between servers and clients, use the following format: <code>ENABLE_DIRECT_CONNECTION = TRUE</code>
Equivalent NetBackup web UI host property	No equivalent exists in the host properties.

ENABLE_NBSQLADM option for NetBackup servers and clients

This option enables or disables the `nbsqladm` command. T

Table 2-113 ENABLE_NBSQLADM information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ENABLE_NBSQLADM = value</pre> The default value is 1. This option should appear only once in the configuration file.
Example	On the server or the client, edit the entry as follows to disable the command: <pre>ENABLE_NBSQLADM = 0</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

FAILOVER_RESTORE_MEDIA_SERVERS option for NetBackup servers

This option specifies that an automatic failover media server be used if a server is temporarily inaccessible for a restore. This failover does not require administrator intervention.

Table 2-114 FAILOVER_RESTORE_MEDIA_SERVERS information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-114 FAILOVER_RESTORE_MEDIA_SERVERS information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>FAILOVER_RESTORE_MEDIA_SERVERS = failed_host host1 host2 ... hostN</pre> ■ <i>failed_host</i> is the server that is not operational. ■ <i>host1 ... hostN</i> are the servers that provide failover capabilities. The default is that NetBackup does not perform automatic failover. When automatic failover is necessary for a server, NetBackup searches from left to right through the associated <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> list. It stops when it finds one that is eligible to perform the restore. Note: The configuration file can contain multiple <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> entries and each entry can list multiple servers. However, a NetBackup server can be a <i>failed_host</i> in only one option. After a <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> option is added, stop and restart the NetBackup Request daemon on the primary server where you plan to change the configuration.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Restore failover. See "Restore failover properties" on page 151.

FETCH_MSDP_EKMS_KEY for NetBackup servers and clients

Starting with NetBackup 11.1, the KMIP Encrypt/Decrypt operations are preferably used to communicate with an external KMS server. If the KMS server does not support the Encrypt/Decrypt operations, NetBackup automatically starts using the GetKey operation to communicate with the KMS server.

With the GetKey operation, the key is transferred from KMS server to NetBackup primary server and this may cause security issues.

Set the `FETCH_MSDP_EKMS_KEY` configuration option to '0' or 'False' to disable the transfer of the key from the KMS server to the NetBackup primary server.

Table 2-115

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To disable the transfer of a key from KMS server to NetBackup primary server., use the following format: <pre>FETCH_MSDP_EKMS_KEY = 0</pre>
Equivalent NetBackup web UI property	No equivalent property in the web UI.

FORCE_RESTORE_MEDIA_SERVER option for NetBackup servers

This option forces the restore to go to a specific server, regardless of where the files were backed up.

Table 2-116 FORCE_RESTORE_MEDIA_SERVER information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>FORCE_RESTORE_MEDIA_SERVER = <i>from_host to_host</i></pre> Where <i>from_host</i> is the server that performed the original backup and <i>to_host</i> is the server to use for the restore. Stop and restart the NetBackup Request daemon on the primary server after adding the <code>FORCE_RESTORE_MEDIA_SERVER</code> option. Physically move the media to <i>to_host</i> before attempting a restore. Update the Media Manager volume database to reflect the move. This setting applies to all storage units on the original server. Restores for any storage unit on <i>from_host</i> go to <i>to_host</i>. To revert to the original configuration for future restores, delete the option.

Table 2-116 FORCE_RESTORE_MEDIA_SERVER information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > General server > Media host override. See “General server properties” on page 106.

GENERATE_ENGLISH_LOGS option for NetBackup servers and clients

This option enables the generation of an English error log, and English trace logs for the `bparchive`, `bpbackup`, `bpduplicate`, `bpimport`, and `bprestore` commands. This option is useful to support personnel to assist in distributed environments where different locales result in logs that contain various languages.

An English text error log (indicated by the suffix `_en`) is created in the following directory:

- On Windows: `Install_path\NetBackup\db\error`
- On UNIX: `/usr/opensv/netbackup/db/error`

Table 2-117 GENERATE_ENGLISH_LOGS information

Usage	Description
Where to use	On NetBackup primary servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GENERATE_ENGLISH_LOGS</pre> This entry should appear only once in the configuration file. Setting the <code>GENERATE_ENGLISH_LOGS</code> option also forces the <code>-en</code> argument on the execution of the following commands when the progress log is specified (<code>-L</code>): <code>bparchive</code>, <code>bpbackup</code>, <code>bpduplicate</code>, <code>bpimport</code>, and <code>bprestore</code>. The suffix <code>_en</code> indicates the English text progress log.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

GUI_ACCOUNT_LOCKOUT_DURATION option for NetBackup servers

This setting determines the amount of time that a user account is locked out after the user exceeds the maximum of failed logon attempts. After that time period the account is unlocked.

Table 2-118 GUI_ACCOUNT_LOCKOUT_DURATION information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GUI_ACCOUNT_LOCKOUT_DURATION = minutes</pre> The default value is 15 minutes.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > User account settings > Unlock locked accounts after. Security > User sessions > User account settings > Unlock locked accounts after.

GUI_IDLE_TIMEOUT option for NetBackup servers

This setting logs out the user session if there is no GUI activity for the specified period of time.

Table 2-119 GUI_IDLE_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-119 GUI_IDLE_TIMEOUT information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GUI_IDLE_TIMEOUT = <i>minutes</i></pre> <code>GUI_IDLE_TIMEOUT</code> is disabled by default.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > User account settings > Session idle timeout. Security > User sessions > User account settings > Session idle timeout.

GUI_MAX_CONCURRENT_SESSIONS option for NetBackup servers

This setting limits the number of concurrent API sessions that a user can have active. API sessions are used for some applications in the NetBackup Administration Console. This setting does not apply to API key sessions or to other applications like the NetBackup Backup, Archive, and Restore interface.

Table 2-120 GUI_MAX_CONCURRENT_SESSIONS information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GUI_MAX_CONCURRENT_SESSIONS = <i>number of sessions</i></pre> Where <i>number of sessions</i> is the number of sessions that users can have open concurrently. <code>GUI_MAX_CONCURRENT_SESSIONS</code> is disabled by default.

Table 2-120 GUI_MAX_CONCURRENT_SESSIONS information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > User account settings > Maximum concurrent sessions. Security > User sessions > User account settings > Maximum concurrent sessions.

GUI_MAX_LOGIN_ATTEMPTS option for NetBackup servers

This setting determines the number of failed logon attempts after which to lock an account out of the NetBackup web UI.

Table 2-121 GUI_MAX_LOGIN_ATTEMPTS information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>GUI_MAX_LOGIN_ATTEMPTS = number of attempts</pre> Where <i>number of attempts</i> is the number of logon attempts after which to lock the user account. The default value is 5. This option is disabled by default.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > User account settings > Number of failed sign-in attempts allowed. Security > User sessions > User account settings > Number of failed sign-in attempts allowed.

HOSTDB_RESYNC_INTERVAL option for NetBackup servers and clients

The `HOSTDB_RESYNC_INTERVAL` option specifies the time interval to synchronize host's information to the NetBackup primary server's host database.

Table 2-122 HOSTDB_RESYNC_INTERVAL information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>HOSTDB_RESYNC_INTERVAL = hours</pre> The default value is 24 hours (1 day). The maximum value for this option is 168 hours (7 days). The minimum value for this option is zero. Setting the option to zero disables periodic updates to the host database. It also disables updates to the host database after the <code>bpcd</code> service restarts. This option should appear only once in the configuration file.
Example	The following example sets the time interval of 72 hours (3 days) to allow data synchronization with the host database: <pre>HOSTDB_RESYNC_INTERVAL = 72</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

HYPERV_WMI_CREATE_DISK_TIMEOUT option for NetBackup servers

This option specifies the timeout period for creating a virtual disk during restore of a Hyper-V VM that was backed up with the WMI method.

Table 2-123 HYPERV_WMI_CREATE_DISK_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-123 HYPERV_WMI_CREATE_DISK_TIMEOUT information
(continued)

Usage	Description
How to use	Change the <code>HYPERV_WMI_CREATE_DISK_TIMEOUT</code> by using the <code>bpsetconfig</code> command or the <code>nbsetconfig</code> command. For information about these commands, see the NetBackup Commands Reference Guide. The default is 24 hours. The range for this option is 0 hours to 240 hours. A value of 0 means the restore job never times out during virtual disk creation.
Example	The following entry tells the NetBackup restore job to wait 48 hours for creation of the Hyper-V virtual disk. <pre>HYPERV_WMI_CREATE_DISK_TIMEOUT = 48</pre> More information on this configuration option is available in the <i>NetBackup for Hyper-V Administrator's Guide</i>.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

INCOMPLETE_JOB_CLEAN_INTERVAL option for NetBackup servers and clients

This option indicates the number of days a failed restore job can remain in the incomplete state before it is moved to the done state.

Table 2-124 INCOMPLETE_JOB_CLEAN_INTERVAL information

Usage	Description
Where to use	On NetBackup primary servers or clients.

Table 2-124 INCOMPLETE_JOB_CLEAN_INTERVAL information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>INCOMPLETE_JOB_CLEAN_INTERVAL = number_of_days</pre> The default is 7 days. Where <i>x</i> is a value between 0 and 365. A value of 0 indicates that failed, incomplete jobs are never automatically moved to the done state. This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or client. > Cleanup > Move restore job from incomplete state to done state. See “Clean up properties” on page 61.

INITIAL_BROWSE_SEARCH_LIMIT option for NetBackup servers and clients

This option specifies the number of days back that NetBackup searches for files to restore. It can improve performance when large numbers of backups are performed.

Table 2-125 INITIAL_BROWSE_SEARCH_LIMIT information

Usage	Description
Where to use	On NetBackup primary servers or clients.

Table 2-125 INITIAL_BROWSE_SEARCH_LIMIT information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>INITIAL_BROWSE_SEARCH_LIMIT = number_of_days</pre> The default is that NetBackup includes files from the time of the last full backup through the latest backup for the client. If the client belongs to more than one policy the browse starts with the earliest of the set of last full backups. This entry should appear only once in the configuration file. When this option is specified on a UNIX client, it applies only to that client. The option can reduce the size of the Search window from what is specified on the server (the client setting cannot make the window larger).
Example	The following example limits the browse range to the seven days before the current date. <pre>INITIAL_BROWSE_SEARCH_LIMIT = 7</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or client > Universal settings > Browse timeframe for restores. See "Universal settings properties" on page 178.

INITIATE_REVERSE_CONNECTION for servers

The `INITIATE_REVERSE_CONNECTION` option lets the primary server or the media server publish the messages to the message queue broker during communication with NAT clients or NAT servers (or NAT hosts).

Table 2-126 INITIATE_REVERSE_CONNECTION information

Usage	Description
Where to use	On NetBackup servers.

Table 2-126 INITIATE_REVERSE_CONNECTION information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To initiate a reverse connection with NAT hosts, use the following format: <pre>INITIATE_REVERSE_CONNECTION = TRUE</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

IP_ADDRESS_FAMILY option for NetBackup servers

The `IP_ADDRESS_FAMILY` option indicates whether NetBackup on this host supports only IPv4 or both IPv4 and IPv6.

If any of the `SERVER` entries do not support IPv4, NetBackup uses the setting that indicates both IPv4 and IPv6.

Table 2-127 IP_ADDRESS_FAMILY information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	On the hosts that use both IPv4 and IPv6, use this option to indicate which address family to use. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>IP_ADDRESS_FAMILY = AF_INET AF_INET6 AF_UNSPEC</pre> <code>AF_INET</code> indicates that the host supports only IPv4. <code>AF_INET6</code> indicates that the host supports only IPv6. <code>AF_UNSPEC</code> indicates that the host supports both IPv4 and IPv6. This entry should appear only once in the configuration file.

Table 2-127 IP_ADDRESS_FAMILY information (continued)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the host > Network settings . See “Network settings properties” on page 124.

JOB_PRIORITY option for NetBackup servers and clients

Use this option to set the priority for a job type.

Table 2-128 JOB_PRIORITY information

Usage	Description
Where to use	On NetBackup primary servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>JOB_PRIORITY = P1 P2 P3 P4 P5 P6 P7 P8 P9 P10 P11 P12 P13 P14 P15 P16 P17 P18 P19 P20 P21 P22 P23 P24</pre> Variables <i>P1</i>, <i>P2</i>, and so on indicate the priority for each backup type. Table 2-129 shows the default priority values. The actual default values for the option appear as follows: <pre>JOB_PRIORITY = 0 0 90000 90000 90000 90000 85000 85000 80000 80000 80000 80000 75000 75000 70000 70000 50000 50000 0 0 0 0 0 0</pre> This entry should appear only once in the configuration file.
Example	To give backup jobs a priority of 50000 and duplication jobs a priority of 30000, change the <code>JOB_PRIORITY</code> entry as follows: <pre>JOB_PRIORITY = 50000 0 90000 90000 90000 90000 85000 85000 80000 80000 80000 80000 75000 75000 70000 70000 30000 50000 0 0 0 0 0 0</pre>

Table 2-128 JOB_PRIORITY information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or client > Default job priorities > Job priority. See “Default job priorities properties” on page 85.

[Table 2-129](#) lists the order of the job types and the various job type defaults.

Table 2-129 Default job type priorities

Field	Represents this action	Default
P1	Performing a backup	0
P2	Performing a database backup (a catalog backup)	0
P3	Performing a restore	90000
P4	Recovering a catalog	90000
P5	Performing a staging operation	90000
P6	Performing the duplication jobs that Vault starts	90000
P7	Cleaning up images	85000
P8	Importing images	85000
P9	Requesting tapes	80000
P10	Cleaning a tape	80000
P11	Tape formatting	80000
P12	Performing device diagnostics	80000
P13	Verifying an image	75000
P14	Running a media contents report	75000
P15	Labeling tape media	70000
P16	Erasing media	70000
P17	Running a duplication job	50000
P18	Performing an inventory	50000
P19	This field is not currently in use	0

Table 2-129 Default job type priorities *(continued)*

Field	Represents this action	Default
P20	This field is not currently in use	0
P21	This field is not currently in use	0
P22	This field is not currently in use	0
P23	This field is not currently in use	0
P24	This field is not currently in use	0

KEEP_LOGS_SIZE_GB for NetBackup servers and clients

The `KEEP_LOGS_SIZE_GB` option specifies the size of the NetBackup logs that you want to retain. When the NetBackup log size grows up to this configuration value, the older logs are deleted.

Table 2-130 `KEEP_LOGS_SIZE_GB` information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 2-130 KEEP_LOGS_SIZE_GB information (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Initially, the <code>bp.conf</code> file does not contain <code>KEEP_LOGS_SIZE_GB</code> entry. Enable the Keep logs up to GB option in the Logging dialog box on the NetBackup Administration Console to set the log retention in GB. A corresponding entry is added in the <code>bp.conf</code> file as follows: <pre>KEEP_LOGS_SIZE_GB = 25</pre> If you disable the Keep logs up to GB option, the <code>bp.conf</code> file shows the corresponding entry as follows: <pre>KEEP_LOGS_SIZE_GB = 0</pre> To set it to a different value, update the <code>bp.conf</code> file using the <code>nbsetconfig</code> command. Use the following format to set <code>KEEP_LOGS_SIZE_GB</code> to a new value in the <code>bp.conf</code> file: <pre>KEEP_LOGS_SIZE_GB = X</pre> 'X' indicates the log size in GB. Note: For NetBackup servers, the recommended value for the <code>KEEP_LOGS_SIZE_GB</code> option is 25 GB or greater. For NetBackup clients, the recommended value for the <code>KEEP_LOGS_SIZE_GB</code> option is 5 GB or greater. This entry should appear only once in the <code>bp.conf</code> configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Logging > Maximum log size. See “Logging properties” on page 113.

KMS_CONFIG_IN_CATALOG_BKUP for NetBackup primary server

Use the `KMS_CONFIG_IN_CATALOG_BKUP` option to include the KMS configuration as part of the disaster recovery (DR) package during catalog backup.

Table 2-131

Usage	Description
Where to use	On NetBackup primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. By default, the <code>KMS_CONFIG_IN_CATALOG_BKUP</code> option is set to '0' (zero). To include the KMS configuration in catalog backup as part of the disaster recovery (DR) package during catalog backup, use the following format: <pre>KMS_CONFIG_IN_CATALOG_BKUP = 1</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

LIMIT_BANDWIDTH option for NetBackup servers

This option specifies a limit for the network bandwidth that NetBackup clients use on a network. The actual limiting occurs on the client side of the backup connection. This option limits only backups. Restores are unaffected. The default is that the bandwidth is not limited.

Note: `LIMIT_BANDWIDTH` applies only to IPv4 networks. Use the `THROTTLE_BANDWIDTH` option to limit IPv6 networks.

See [“THROTTLE_BANDWIDTH option for NetBackup servers”](#) on page 315.

See [“Throttle bandwidth properties”](#) on page 174.

MALWARE_DETECTION_JOBS_PER_SCAN_HOST option for NetBackup servers

The `MALWARE_DETECTION_JOBS_PER_SCAN_HOST` parameter is used to configure the number of parallel scans that are allowed on each scan host.

Table 2-132 MALWARE_DETECTION_JOBS_PER_SCAN_HOST option information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: MALWARE_DETECTION_JOBS_PER_SCAN_HOST = 5 By default: ■ The number of parallel scans per scan host limit is 3. ■ The minimum supported value is 1. ■ The maximum supported value is 10.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

MALWARE_SCAN_OPERATION_TIMEOUT

The **MALWARE_SCAN_OPERATION_TIMEOUT** parameter is used to configure the duration of the scan operation that is allowed to run before timeout happens.

Scan operation for backup image can take a long time based upon the factors like backup size, number of files in the backup. By default, scan operation times out after 2 days. User can set the timeout value from 1 hour to 30 days.

Table 2-133 MALWARE_SCAN_OPERATION_TIMEOUT option information

Usage	Description
Where to use	On NetBackup media servers.

Table 2-133 MALWARE_SCAN_OPERATION_TIMEOUT option information
(continued)

Usage	Description
How to use	Use <code>nbgetconfig</code> or <code>nbsetconfig</code> commands to view, add, or change the value of the timeout. For information about these commands, see the NetBackup Commands Reference Guide. Set the configuration key on the MSDP media server where <code>ScanManager</code> (<code>nbcs</code>) is started. For multiple MSDP media servers, set the configuration key on each server. Use the following format: MALWARE_SCAN_OPERATION_TIMEOUT = 120 By default scan operation timeout value is 2880 minutes (2 days). The minimum supported value is 60 minutes (1 hour) and the maximum supported value is 43200 minutes (30 days).
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

MANAGE_WIN_CERT_STORE_PRIVATE_KEY option for NetBackup primary servers

The `MANAGE_WIN_CERT_STORE_PRIVATE_KEY` option lets you disable the automatic permission management of the private key of the certificate in Windows Certificate Store.

This option is applicable for Windows Certificate Store and only when the NetBackup services are running in the Local Service account context.

When NetBackup services are running in the Local Service account context, the services need to have permissions to read the private key for certificate in Windows Certificate Store.

When the `MANAGE_WIN_CERT_STORE_PRIVATE_KEY` option is set to `Automatic`, the NetBackup service that is running in the privileged user account context grants access to all other NetBackup services for reading the private key whenever required.

By default, permissions for the private key are automatically managed. When the `MANAGE_WIN_CERT_STORE_PRIVATE_KEY` option is set to `Disabled`, the permissions of the private key need to be managed manually.

Note: It is not recommended to set the `MANAGE_WIN_CERT_STORE_PRIVATE_KEY` option to `Disabled`.

To manually update the permissions when this option is `Disabled`, run the following command:

```
nbccertcmd -setWinCertPrivKeyPermissions -reason audit reason -force
```

Refer to the [NetBackup Commands Reference Guide](#) for more details on the command-line options.

Table 2-134 `MANAGE_WIN_CERT_STORE_PRIVATE_KEY` information

Usage	Description
Where to use	On NetBackup primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MANAGE_WIN_CERT_STORE_PRIVATE_KEY = Automatic</pre>
Equivalent NetBackup web UI property	No equivalent exists.

MAX_LOGFILE_SIZE option for NetBackup servers and clients for legacy logging

The `MAX_LOGFILE_SIZE` option specifies the maximum size that you want to set for a log file. When the log file size in NetBackup matches the `MAX_LOGFILE_SIZE` setting, the next logs are stored in a new log file.

Table 2-135 `MAX_LOGFILE_SIZE` information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 2-135 MAX_LOGFILE_SIZE information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> (or <code>bpgetconfig</code>) and the <code>nbsetconfig</code> (or <code>bpsetconfig</code>) commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format to set <code>MAX_LOGFILE_SIZE</code> to a new value in the <code>bp.conf</code> file: <pre>MAX_LOGFILE_SIZE = X</pre> 'X' indicates maximum size of a NetBackup log file in MB. Note: <code>MAX_LOGFILE_SIZE</code> should be set to an integer number, which should be greater than '0'. If you have set <code>MAX_LOGFILE_SIZE</code> to an invalid value such as 0 or -100, it is automatically set to the default value (500 MB). This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

MAX_NUM_LOGFILES option for NetBackup servers and clients for legacy logging

The `MAX_NUM_LOGFILES` option specifies the maximum number of log files that you want to be retained in a NetBackup log directory. When the number of log files in the directory matches the `MAX_NUM_LOGFILES` setting, the oldest log file is deleted.

Table 2-136 MAX_NUM_LOGFILES information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 2-136 MAX_NUM_LOGFILES information (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> (or <code>bpgetconfig</code>) and the <code>nbsetconfig</code> (or <code>bpsetconfig</code>) commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format to set <code>MAX_NUM_LOGFILES</code> to a new value in the <code>bp.conf</code> file: <pre>MAX_NUM_LOGFILES = X</pre> 'X' indicates maximum number of NetBackup log files that are created in a log directory. <code>MAX_NUM_LOGFILES</code> should be set to a number that should be greater than one (1). If you have set <code>MAX_NUM_LOGFILES</code> to an invalid value such as 0 or 1, it is automatically set to the default value, which is infinite. However, in the <code>bp.conf</code> file, the <code>MAX_NUM_LOGFILES</code> entry appears as follows: <pre>MAX_NUM_LOGFILES = 0</pre> Zero (0) indicates an infinite value. This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

MEDIA_UNMOUNT_DELAY option for NetBackup servers

When `MEDIA_UNMOUNT_DELAY` is specified, the media unload is delayed for the specified number of seconds after the requested operation has completed. (Applies only to user operations.)

Table 2-137 MEDIA_UNMOUNT_DELAY information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-137 MEDIA_UNMOUNT_DELAY information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MEDIA_UNMOUNT_DELAY = seconds</pre> The default is a media mount delay time of 180 seconds. This entry should appear only once in the configuration file.
Example	The delay is set to 120 seconds in the following example: <pre>MEDIA_UNMOUNT_DELAY = 120</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Media > Media unmount delay. See “Media properties” on page 119.

MEDIA_REQUEST_DELAY option for NetBackup servers

This option specifies the number of seconds that NetBackup waits for a non-robotic drive to become ready.

Table 2-138 MEDIA_REQUEST_DELAY information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Change the <code>MEDIA_REQUEST_DELAY</code> by using the <code>nbemmcmd</code> command. For information about these commands, see the NetBackup Commands Reference Guide. The default is that NetBackup does not wait for non-robotic drives to become ready. This option should appear only once in the configuration file.
Example	The following command indicates to NetBackup to wait 150 seconds for a non-robotic drive to become ready for use. <pre>nbemmcmd -changesetting -MEDIA_REQUEST_DELAY 150</pre>

Table 2-138 MEDIA_REQUEST_DELAY information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Media > Media request delay . See “Media properties” on page 119.

MEDIA_SERVER option for NetBackup servers

This option is similar to the `SERVER` option.

A host that is listed as a `MEDIA_SERVER` can back up and restore clients. However, if the host is not specified as a `SERVER`, the host has limited administrative capabilities.

Table 2-139 MEDIA_SERVER information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MEDIA_SERVER = media_server_name</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Servers > Media servers . See “Servers properties” on page 159.

MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB option for NetBackup servers

This high water mark for resilient backups specifies the amount of free space to maintain on the disk. Backup jobs waiting to reestablish communications with the primary server pause if they cannot cache metadata. The value is specified in megabytes. The minimum value is zero (0), which means use all available disk space.

Be aware of the relationship between the `RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH` and `MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB` values. Use `RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH` to specify where to write the

cache information. Use `MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB` to specify how much free disk space to maintain.

Table 2-140

Usage	Description
Where to use	On NetBackup media servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB = number</pre> This option should appear only once in the configuration file. The default value for <code>MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB</code> is 5000 MB.
Example	On the media server, add the following entry to specify the amount of free space to maintain on the disk: <pre>MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB = 100</pre>

MPX_RESTORE_DELAY option for NetBackup servers

This option applies to multiplexed restores. The `MPX_RESTORE_DELAY` specifies how long the server waits for restore requests of files and raw partitions. In this case, the option applies to the files and raw partitions in a set of multiplexed images on the same tape. All of the restore requests that are received within the delay period are included in the same restore operation (one pass of the tape).

Table 2-141 MPX_RESTORE_DELAY information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-141 MPX_RESTORE_DELAY information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MPX_RESTORE_DELAY = seconds</pre> The default is 30 seconds. This option should appear only once in the configuration file.
Example	The following example configures a server to wait 15 minutes. <pre>MPX_RESTORE_DELAY = 900</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > General server > Delay on multiplexed restores. See “General server properties” on page 106.

MUST_USE_LOCAL_DRIVE option for NetBackup servers

This option instructs NetBackup that if the client is also a media server and this option is present, backups for this client must occur on a local drive. If all drives are down, another may be used. If the client is not a media server, this option has no effect.

Table 2-142 MUST_USE_LOCAL_DRIVE information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbemmcmd</code> command to change the option. For example: <pre>nbemmcmd -changesetting -MUST_USE_LOCAL_DRIVE yes</pre> For information about these commands, see the NetBackup Commands Reference Guide. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > General server > Must use local drive. See “General server properties” on page 106.

NAT_SERVER_LIST for servers

The `NAT_SERVER_LIST` option is used to specify NAT servers with which NetBackup servers in a public network can establish a reverse connection. If the option is not configured, the NetBackup servers are considered to be in the same network.

Table 2-143 NAT_SERVER_LIST information

Usage	Description
Where to use	On NetBackup servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>NAT_SERVER_LIST = NAT server 1 NAT server 2</pre> The NAT server names should be separated by spaces.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NB_FIPS_MODE option for NetBackup servers and clients

Use the `NB_FIPS_MODE` option to enable the FIPS mode in your NetBackup domain.

Table 2-144 NB_FIPS_MODE information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. By default, the <code>NB_FIPS_MODE</code> option is disabled. To enable the option, use the following format: <pre>NB_FIPS_MODE = ENABLE</pre> To disable the option, use the following format: <pre>NB_FIPS_MODE = DISABLE</pre>

Table 2-144 NB_FIPS_MODE information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBFP_CHECK option for NetBackup servers and clients

Use the `NBFP_CHECK` option to control whether permissions and ownership of NetBackup files should be verified, corrected, and reported to the web UI or not.

Table 2-145

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. The <code>NBFP_CHECK</code> configuration option can have the following values: <code>FIX</code> – Corrects permissions and ownership of the files. <code>FIX_NOTIFY</code> – Corrects permissions and ownership of the files, and posts the results on to the NetBackup web UI. <code>VERIFY</code> – Verifies permissions and ownership of the files. <code>VERIFY_NOTIFY</code> – Verifies permissions and ownership, and posts the results on to the NetBackup web UI. To disable the option, use the following format: <code>NBFP_CHECK = OFF</code>
Equivalent NetBackup web UI property	No equivalent exists in the NetBackup web UI.

NBRNTD_IDLE_TIMEOUT option for NetBackup servers

The `NBRNTD_IDLE_TIMEOUT` option specifies the number of seconds to wait before the Remote Network Transport Service (`nbrntd`) shuts itself down. The Remote Network Transport Service provides resilient network connections. After it is shut down, NetBackup must restart the service if a new resilient connection for backups or restores is required.

Table 2-146 NBRNTD_IDLE_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>NBRNTD_IDLE_TIMEOUT = <i>seconds</i></pre> The default timeout is 300 seconds (five minutes). By default, this entry is not present in the configuration file.
Example	In the following example, the Remote Network Transport Service shuts off after 15 minutes. <pre>NBRNTD_IDLE_TIMEOUT = 900</pre>

See “[RESILIENT_NETWORK option for NetBackup primary servers and clients](#)” on page 302.

NBSD_POLL_INTERVAL option for NetBackup servers and clients

The `NBSD_POLL_INTERVAL` option specifies the interval in seconds after which the service checks the status of the registered process. The default value is 600.

Table 2-147 NBSD_POLL_INTERVAL Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-147 NBSD_POLL_INTERVAL Information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>NBSD_POLL_INTERVAL = <i>seconds</i></pre> Example: <pre>NBSD_POLL_INTERVAL = 20</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_DUMP_COUNT option for NetBackup servers and clients

The `NBSD_DUMP_COUNT` option specifies the maximum number of process dumps that are collected for a registered process. You can specify the value between 1 and 10. Default value is 3.

Table 2-148 NBSD_DUMP_COUNT Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>NBSD_DUMP_COUNT = <i>numbers</i></pre> Example: <pre>NBSD_DUMP_COUNT = 3</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_MONITOR_CPU option for NetBackup servers and clients

The `NBSD_MONITOR_CPU` option specifies the process to monitor the CPU usage in percentage.

Note: Do not specify percentage in (%) sign.

Table 2-149 NBSD_MONITOR_CPU Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>PROC_NAME1:CPU_percent, PROC_NAME2:CPU_percent, PROC_NAME3:CPU_percent</pre> <pre>NBSD_MONITOR_CPU = bpdbrm:40, bpbbrm:50</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_MONITOR_MEMORY option for NetBackup servers and clients

The `NBSD_MONITOR_MEMORY` option specifies the process to monitor the memory usage.

Note: Do not enclose the value in single or double quotes.

Table 2-150 NBSD_MONITOR_MEMORY Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-150 NBSD_MONITOR_MEMORY Information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>PROC_NAME1:MEM_SIZE1, PROC_NAME2:MEM_SIZE2, PROC_NAME3:MEM_SIZE3</pre> <pre>NBSD_MONITOR_MEMORY = nbsl:8196, bpdbrm:4096</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_MEMORY_UNIT option for NetBackup servers and clients

The `NBSD_MEMORY_UNIT` option specifies the unit for the memory calculations which is used in the process from the total system memory. If the unit is `PERCENT`, then the calculations are based on the percent of memory. If the unit is `ABSOLUTE`, then the calculations are based on the absolute value in MB. Default value is `ABSOLUTE`.

Note: Do not enclose the value in single or double quotes.

Table 2-151 NBSD_MEMORY_UNIT Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>NBSD_MEMORY_UNIT = ABSOLUTE</pre> By default, this entry is not present in the configuration file.

Table 2-151 NBSD_MEMORY_UNIT Information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_MONITOR_DEADLOCK option for NetBackup servers and clients

The `NBSD_MONITOR_DEADLOCK` option specifies the process to monitor the deadlock. The CPU and memory usage are the frequent long intervals which are assumed as deadlock. Default value is 60 minutes, but you can set the value based on the case.

Note: Do not enclose the value in single or double quotes.

Table 2-152 NBSD_MONITOR_DEADLOCK Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>PROC_NAME1, PROC_NAME2, PROC_NAME3 NBSD_MONITOR_DEADLOCK = nbwmc, adminconsole, nbsl</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_DEADLOCK_INTERVAL option for NetBackup servers and clients

The `NBSD_DEADLOCK_INTERVAL` is an interval after which to consider that the process is inactive. Default value is 60.

Table 2-153 NBSD_DEADLOCK_INTERVAL Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>NBSD_DEADLOCK_INTERVAL = <i>minutes</i></pre> Example: <pre>NBSD_DEADLOCK_INTERVAL = 60</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_ALWAYS_DUMP option for NetBackup servers and clients

The `NBSD_ALWAYS_DUMP` option specifies the service to always dump whenever the CPU breaks the threshold or to let the service manage it logically based on the average of previous readings. Default value is 0.

Table 2-154 NBSD_ALWAYS_DUMP Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>NBSD_ALWAYS_DUMP = 0</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_CAPTURE_PROCESS_DUMP option for NetBackup servers and clients

The `NBSD_CAPTURE_PROCESS_DUMP` option specifies whether to capture the process dump. You can set the value to 1, if you want to capture the process dump after the threshold is reached. Default value is 1.

Table 2-155 NBSD_CAPTURE_PROCESS_DUMP Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>NBSD_CAPTURE_PROCESS_DUMP = 0</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_INCREASE_LOG_LEVEL option for NetBackup servers and clients

The `NBSD_INCREASE_LOG_LEVEL` option specifies to increase the log levels for the VXUL processes automatically and make changes in the `nblog.conf` file. Logs for the legacy processes are not changed as it might lead to huge levels.

Table 2-156 NBSD_INCREASE_LOG_LEVEL Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-156 NBSD_INCREASE_LOG_LEVEL Information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_INCREASE_LOG_LEVEL = 1</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_CAPTURE_NETWORK_STAT option for NetBackup servers and clients

The `NBSD_CAPTURE_NETWORK_STAT` option specifies to set the value as 1, if you want to monitor the network connections at the time of an event. Default value is 1.

Table 2-157 NBSD_CAPTURE_NETWORK_STAT Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_CAPTURE_NETWORK_STAT = 1</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_CAPTURE_DISK_IO option for NetBackup servers and clients

The `NBSD_CAPTURE_DISK_IO` option specifies to set the value as 1, if you want to capture the system DISK IO stats at the point of an event. Default value is 1.

Table 2-158 NBSD_CAPTURE_DISK_IO Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>NBSD_CAPTURE_DISK_IO = 1</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_NUMBER_OF_READINGS option for NetBackup servers and clients

The `NBSD_NUMBER_OF_READINGS` option specifies the number of reading to capture in case of the threshold event. Default value is 50.

Table 2-159 NBSD_NUMBER_OF_READINGS Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>NBSD_NUMBER_OF_READINGS = 50</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_READING_INTERVAL option for NetBackup servers and clients

The `NBSD_READING_INTERVAL` option specifies to take the reading at a specific interval. Default value is 5.

Table 2-160 NBSD_READING_INTERVAL Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <pre>NBSD_READING_INTERVAL = 5</pre> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_PURGE_OLD_EVIDENCE option for NetBackup servers and clients

The `NBSD_PURGE_OLD_EVIDENCE` option purges the old evidences present in the directory given in `NBSD_EVIDENCE_PATH`. You must copy the old evidence before it is lost. Default value is 0.

Table 2-161 NBSD_PURGE_OLD_EVIDENCE Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-161 NBSD_PURGE_OLD_EVIDENCE Information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: NBSD_PURGE_OLD_EVIDENCE = 0 By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_CAPTURE_WITHOUT_THRESHOLD option for NetBackup servers and clients

The `NBSD_CAPTURE_WITHOUT_THRESHOLD` option specifies to capture the evidence set without any threshold for the registered processes. This takes precedence over all the threshold flags. It is required in case we need to gather the evidence without any threshold event but at a set regular interval. Default value is 0.

Table 2-162 NBSD_CAPTURE_WITHOUT_THRESHOLD Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: NBSD_CAPTURE_WITHOUT_THRESHOLD = 1 By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_JDK_HOME option for NetBackup servers and clients

The `NBSD_JDK_HOME` option specifies the path to the JDK Home folder. The path is required to execute JSTACK on a java process. For example: If JDK is installed in `c:\jdk`, `JAVA_HOME` should be set to `c:\jdk`.

Note: Do not enclose the value in single or double quotes.

Table 2-163 NBSD_JDK_HOME Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_JDK_HOME = c:\jdk1.8</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_EVIDENCE_PATH option for NetBackup servers and clients

The `NBSD_EVIDENCE_PATH` option specifies the path to the folder where you want to store the generated evidence. It is a mandatory value. The `nbperfmon` does not run if the value is not specified. Ensure, the folder has sufficient space to gather the logs.

Note: Do not enclose the value in single or double quotes.

Table 2-164 NBSD_EVIDENCE_PATH Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-164 NBSD_EVIDENCE_PATH Information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: NBSD_EVIDENCE_PATH = c:\\temp By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_VERBOSE option for NetBackup servers and clients

The `NBSD_VERBOSE` option specifies to enable the verbose logs for the performance of the NetBackup Smart Diagnosis (NBSD). Default value is 0.

Table 2-165 NBSD_VERBOSE Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: NBSD_VERBOSE = 1 By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_AUTO_MONITOR option for NetBackup servers and clients

The `NBSD_AUTO_MONITOR` option enables to monitor all the NetBackup processes on the NetBackup host with default CPU and Memory threshold values. The default CPU threshold is 90% and default memory threshold are 60%.

To change default threshold values refer
NBSD_AUTOMONITOR_CPU_THRESHOLD and
NBSD_AUTOMONITOR_MEMORY_THRESHOLD configuration parameters.

Note: NBSD_CAPTURE_WITHOUT_THRESHOLD parameter is ignored when
NBSD_AUTO_MONITOR is set to 1.

Table 2-166 NBSD_AUTO_MONITOR Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: NBSD_AUTO_MONITOR = 0 By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_AUTOMONITOR_CPU_THRESHOLD option for NetBackup servers and clients

The `NBSD_AUTOMONITOR_CPU_THRESHOLD` option enables to monitor all the NetBackup processes in the NetBackup host with default CPU and Memory threshold values. The value is in percentage and can have values in between 1 and 100.

Table 2-167 NBSD_AUTOMONITOR_CPU_THRESHOLD Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-167 NBSD_AUTOMONITOR_CPU_THRESHOLD Information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_AUTOMONITOR_CPU_THRESHOLD = 90</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_AUTOMONITOR_MEMORY_THRESHOLD option for NetBackup servers and clients

The `NBSD_AUTOMONITOR_MEMORY_THRESHOLD` option defines the memory threshold value used for monitoring all the NetBackup processes. The value denotes the memory usage percentage of the total memory process. This value is considered for threshold decision making. The value is in percentage and can have values between 1 and 100.

Table 2-168 NBSD_AUTOMONITOR_MEMORY_THRESHOLD Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_AUTOMONITOR_MEMORY_THRESHOLD = 60</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_MONITOR_POLICY_NAME option for primary server

The `NBSD_MONITOR_POLICY_NAME` option specifies the policies to monitor the CPU, memory, and deadlock thresholds. The process launched after the policy execution is automatically added for monitoring with default threshold values.

This parameter is only applicable for primary server. The value for this parameter is a comma separated list of polices to monitor.

`NBSD_MONITOR_POLICY_NAME=Policy1, Policy2, Policy3.`

Note: Do not enclose the value in quotation marks.

The default CPU threshold is 90 percent and default memory threshold are 60%.

To change default threshold values refer `NBSD_AUTOMONITOR_CPU_THRESHOLD` and `NBSD_AUTOMONITOR_MEMORY_THRESHOLD` configuration parameters.

Note: Only the clients mentioned in the policy are considered for monitoring and the processes that are launched on those policy clients after the policy execution are monitored.

`NBSD_CAPTURE_WITHOUT_THRESHOLD` parameter is ignored when `NBSD_MONITOR_POLICY_NAME` is set.

Table 2-169 `NBSD_MONITOR_POLICY_NAME` Information

Usage	Description
Where to use	Only on the primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_MONITOR_POLICY_NAME = Policy1</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_MONITOR_SYSTEM_FOR_HOURS option for NetBackup servers and clients

The `NBSD_MONITOR_SYSTEM_FOR_HOURS` option defines the time in hours after which the `nbsmartdiag` process automatically stops. By default, the service runs for 7 days (168 hrs) and then stops itself. The value 0 means the process runs forever.

Table 2-170 NBSD_MONITOR_SYSTEM_FOR_HOURS Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_MONITOR_SYSTEM_FOR_HOURS = 168</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_EVIDENCE_SIZE_LIMIT option for NetBackup servers and clients

The `NBSD_EVIDENCE_SIZE_LIMIT` option defines the size limit in GB in the evidence folder. The value of 0 means no limitation on size.

Note: If the size before the evidence captured is less than the set size, the evidence is captured and not stopped in between even if the size exceeds. The next evidence is not captured.

Table 2-171 NBSD_EVIDENCE_SIZE_LIMIT Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.

Table 2-171 NBSD_EVIDENCE_SIZE_LIMIT Information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_EVIDENCE_SIZE_LIMIT = 0</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NBSD_PUSH_MONITOR_DATA_TO_REMOTE option for NetBackup servers and clients

The `NBSD_PUSH_MONITOR_DATA_TO_REMOTE` option Allows `nbsmartdiag` to push the list of identified processes with the default threshold values during the policy execution to the respective clients or media servers.

This option is effective when a user mentions the policy name into the `NBSD_MONITOR_POLICY_NAME`.

Table 2-172 NBSD_PUSH_MONITOR_DATA_TO_REMOTE Information

Usage	Description
Where to use	On NetBackup primary, media, or client servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Example: <code>NBSD_PUSH_MONITOR_DATA_TO_REMOTE = 0</code> By default, this entry is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NETBACKUP_NATIVE_AUDITING option for NetBackup primary server

Use the `NETBACKUP_NATIVE_AUDITING` option to stop storing the NetBackup audit events in the NetBackup database.

By default, the `NETBACKUP_NATIVE_AUDITING` option is enabled and the NetBackup audit events are stored in the NetBackup database. If you disable this option, the `nbauditreport` command, the **NetBackup Administration Console** and the NetBackup web UI do not display any audit events.

Disabling the `NETBACKUP_NATIVE_AUDITING` option is not recommended.

Table 2-173 `NETBACKUP_NATIVE_AUDITING` information

Usage	Description
Where to use	On NetBackup primary server.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To stop storing the NetBackup audit events in the NetBackup database, use the following format: <pre>NETBACKUP_NATIVE_AUDITING = 0</pre> This setting is not recommended.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

NOTIFY_SNOOZE_PERIOD_IN_DAYS option for NetBackup servers

The NetBackup administrator may temporarily require certain files such as `No.Restrictions` or `NOexpire` to carry out specific tasks in the NetBackup domain.

You must remove such files after the required tasks are complete. If these files are not removed, notifications are generated and shown in the NetBackup web UI every 7 days.

You can use the `NOTIFY_SNOOZE_PERIOD_IN_DAYS` option to change the frequency of these notifications from the default value to any value from 1 to 90.

No.Restrictions file

The NetBackup administrator can allow clients to perform redirected restores. With the `No.Restrictions` file approach, all NetBackup clients can restore the backups that belong to other clients.

See [“Allowing all clients to perform redirected restores”](#) on page 1115.

Caution: For security reasons, it is strongly advised not to use the `No.Restrictions` file approach.

NOexpire file

The NetBackup administrator may require to use `NOexpire` file if automatic image expiration needs to be temporarily halted in ceratin scenarios, for example during maintenance window.

Caution: For security reasons, it is strongly advised that you remove the `NOexpire` file from the primary server to resume automatic backup image expiration in the image database.

Table 2-174 NOTIFY_SNOOZE_PERIOD_IN_DAYS option

Usage	Description
Where to use	On NetBackup servers
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view the option. For information about these commands, see the <i>NetBackup Commands Reference Guide</i>: http://www.veritas.com/docs/DOC5332 For information about these commands, see the NetBackup Commands Reference Guide. The default value of the option is 7. You can change the value and set it to a value from 1 to 90. Use the following format: <pre>NOTIFY_SNOOZE_PERIOD_IN_DAYS = Number of days</pre> The following example configures a snooze period of 5 days. <pre>NOTIFY_SNOOZE_PERIOD_IN_DAYS = 5</pre>
Equivalent UI property	No equivalent UI property exists.

PREFERRED_NETWORK option for NetBackup servers

The `PREFERRED_NETWORK` option is not needed in an environment if NetBackup is configured with appropriate host names. The operating system must resolve to the correct IP addresses and then route the addresses correctly.

When external constraints prevent the environment from being corrected, `PREFERRED_NETWORK` entries can be useful in the following situations:

- To prevent NetBackup from connecting to specific destination addresses.
- To cause NetBackup to connect only to specific destination addresses.
- To request a subset of local interfaces for source binding when outbound connections are made.

Caution: When used for source binding, the operating system may not honor the source binding list provided by NetBackup. If the operating system implements the weak host model, asymmetrical network routing may result. If asymmetrical routing occurs, the remote host may reject the inbound connection if it implements the strong host model. Similarly, stateful network devices may also drop asymmetrical connections. To ensure the use of specific outbound interfaces for specific remote hosts or networks, make sure that the OS name resolution and routing configurations are correct; create static host routes if needed. Ensure that all network drivers properly implement the IP and TCP networking protocols.

The local `PREFERRED_NETWORK` entries do not affect the forwarding profile that the local host returns to a remote host during initial CORBA connection setup; it will contain all the local plumbed interfaces. However, the End Point Selection algorithm within the remote process will utilize its local `PREFERRED_NETWORK` entries to evaluate the profile when selecting the destination for the subsequent CORBA connection.

Table 2-175 `PREFERRED_NETWORK` information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 2-175 PREFERRED_NETWORK information (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. The option uses the following syntax: <pre>PREFERRED_NETWORK = target[/subnet] directive [source[/subnet]]</pre> Note: The <code>source</code> option is not allowed for the <code>PROHIBITED</code> directive. Multiple <code>PREFERRED_NETWORK</code> entries can be specified. During evaluation, the entries are sorted by length of target subnet. Entries with the largest (more precise) subnet are compared before entries with a shorter (less precise) subnet. If two entries have equal subnet specification, they are compared in the order configured, from the top of the list to the bottom. If a subnet is not specified, the default is <code>/128</code> when the address is non-zero and <code>/0</code> when the address is 0. This applies to both <code>target</code> and <code>source</code> addresses. A subnet of <code>/0</code> cannot be used with a non-zero address because it effectively negates all of the bits in the address, making the <code>target</code> or the <code>source</code> match every address. For example, <code>0/0</code>. The following topics describe details about each option: ■ See “target[/subnet]” on page 286. ■ See “directive” on page 287. ■ See “source[/subnet]” on page 288.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the host > Preferred network. See “Preferred network properties” on page 129.

See [“PREFERRED_NETWORK examples”](#) on page 289.

target[/subnet]

The `target[/subnet]` option indicates a host name or range of addresses to be compared to the prospective source or destination addresses being evaluated. The following are examples of how to indicate a target or a subnet:

A host name	<code>myserver.domain</code>
An IP address	<code>10.82.105.11</code>

A network with subnet	10.82.105.0/21
Any IPv4 address	0.0.0.0
Any IPv6 address	0::0
Any address	0/0

A host or a network name that cannot resolve causes the `target` to be ignored. However, any associated `source` is added to the source binding list.

directive

The `directive` option determines how the `target` is compared to the source and/or the destination address that is being evaluated. The following directives can be used:

MATCH	Applies to destination addresses. If the address that is being evaluated matches the <code>target</code>, then the address is immediately selected to be used and evaluation stops. If the <code>target</code> is not matched, evaluation continues with the next entry.
ONLY	Applies to destination addresses. If the address that is being evaluated does not match the <code>target</code>, it is not used and evaluation stops for this address. If this was the only potential destination, the connection is not attempted. If there is an additional potential destination, it is evaluated starting over with the first entry.
PROHIBITED	The <code>target</code> applies to both source and destination addresses. If a source is specified, it is ignored and the <code>target</code> is prohibited. If the matched address is a destination address, evaluation stops. If this was the only potential destination, the connection is not attempted. If there are additional potential destinations, they are evaluated starting over with the first entry. If the matched address is a source address, it is removed from the binding list. However, if <code>source</code> entries exist, the shortened binding list may not be used. (See “source/subnet” on page 288.) Caution: On some platforms, prohibiting a local interface may cause unexpected results when connecting to remote hosts. Prohibiting a local interface does not affect connections that are internal to the host.

source[/subnet]

source[/subnet] is optional and indicates a host name or IP address range that is requested to be used as the local interface for outbound connection to addresses in the `target`.

source[/subnet] is applicable to the directives `ONLY` and `MATCH`, but not to the directive `PROHIBITED`.

Notes:

- The operating system may not honor the source request.
- NetBackup does not request a *source* which has already been `PROHIBITED`.
- A host or network name that cannot be resolved, or that is not local to the host, is ignored, but the evaluation status of the `target` is still used.

Source binding evaluation

The prospective source binding list is provided by the operating system and consists of all of the local plumbed interfaces that are not loopback (`127.0.0.0/8`, `:::1`) and not link-local (`169.254.0.0/16`, `fe80::/64`).

The `PREFERRED_NETWORK` entries are then evaluated in the sort order by longest subnet first, then position when subnets are equal. Any local interfaces that match a *source* entry are moved to a second list if still present in the first list. Any local interfaces that match a `target PROHIBITED` entry are removed from the first list if not already moved to the second list.

If local interfaces were moved to the second list, that list becomes the tentative binding list. If the second list is empty, and interfaces were removed from the first list, then the shortened first list becomes the tentative binding list.

If a destination matches the `target` of an entry with a *source*, the tentative binding list is searched for the first match to an entry in *source*. If a match is found, that interface becomes the source requested by NetBackup for the outbound connection.

Otherwise, if the tentative binding list is the shortened first list, then it becomes the source binding list for the outbound connection.

Otherwise, `ANY` becomes the source binding list for the outbound connection.

Note: If the source binding list is not `ANY`, most operating systems will arbitrarily use the first interface in the list provided by the application. Because of this behavior, it is best to not use `PROHIBITED` entries for local interfaces and to minimize the use of *source* entries.

PREFERRED_NETWORK examples

Table 2-176 Basic examples

Description	Entry
Allows connectivity to the host names that resolve to 12.34.0.0 through 12.34.255.255. It does not affect outbound interface selection:	PREFERRED_NETWORK = 12.34.0.0/16 MATCH
Allows connectivity to the host name that resolves to 12.34.56.78, and requests that the operating system use 98.76.54.32 as the outbound interface.	PREFERRED_NETWORK = 12.34.56.78 MATCH 98.76.54.32
Instructs the host to use the interface IPs of <i>Host_A</i> for all IPv4 and IPv6 addresses.	PREFERRED_NETWORK = 0/0 MATCH <i>Host_A</i>
Prevents NetBackup from connecting to any destination address outside the range of 12.34.56.0 through 12.34.56.255. The source interface will be ANY unless one or more are PROHIBITED.	PREFERRED_NETWORK = 12.34.56.78/24 ONLY
Prevents NetBackup from connecting to any destination address outside the range of 12.34.56.0 through 12.34.56.255. Requests that the operating system use 98.76.54.32 as the outbound interface.	PREFERRED_NETWORK = 12.34.56.78/24 ONLY 98.76.54.32
Prevents NetBackup from connecting to any destination addresses outside of the indicated IPv6 subnet. The source interface will be ANY unless one or more are PROHIBITED.	PREFERRED_NETWORK = 2001:1234:1efc::/48 ONLY
Prevents NetBackup from using any address between 12.34.0.0 and 12.34.255.255 as the source or destination for a connection. If it matches a local interface, NetBackup will provide an ordered list of the remaining interfaces as the source binding list for the outbound interface when other entries do not specify a <i>source</i> . Using PROHIBITED with local interfaces is discouraged. See the details in the following topic: See " directive " on page 287.	PREFERRED_NETWORK = 12.34.56.78/16 PROHIBITED
Prevents the host from using IPv4 addresses.	PREFERRED_NETWORK = 0.0.0.0 PROHIBITED
Prevents the host from using IPv6 addresses.	PREFERRED_NETWORK = 0::0 PROHIBITED
Prevents the host from using the address of the <i>production_server</i> .	PREFERRED_NETWORK = <i>production_server</i> PROHIBITED

Using `bplocaladdrs` to troubleshoot

Use the `bplocaladdrs` command to observe the local interfaces that are provided to NetBackup by the operating system and the order in which they are provided.

`bplocaladdrs` returns the following output for the host (`bob`) in the examples in the following topics.

```
$ bplocaladdrs
10.82.105.11
10.82.105.8
10.82.10.10
```

Using `bptestnetconn` to troubleshoot

Use the `bptestnetconn` command to observe the order in which entries are evaluated and the evaluation results. The `TGT` or `SRC` indicates whether the destination is permitted and which source binding list NetBackup provides to the operating system. A value of `ANY` indicates that the outbound interface is not constrained by NetBackup.

```
$ bptestnetconn -asp -v6
...
FL: myprimary -> 10.82.105.14 : 5 ms FAST (< 5 sec) TGT PROHIBITED
FL: mymedia -> 10.81.40.61 : 6 ms FAST (< 5 sec) SRC:
10.82.10.10
...
```

`PREFERRED_NETWORK` rules are applied in this order:

```
[0] PREFERRED_NETWORK = 10.82.105.14 PROHIBITED
[1] PREFERRED_NETWORK = 10.81.40.0/24 MATCH 10.82.10.10
```

```
$ bptestnetconn -asp -v6 -H myclient
...
FL: myclient -> 10.81.40.127 : 6 ms FAST (< 5 sec) SRC: ANY
```

`PREFERRED_NETWORK` rules are applied in this order:

```
[0] PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.0/24
[1] PREFERRED_NETWORK = 10.82.105.0/29 PROHIBITED
[2] PREFERRED_NETWORK = 10.82.104.0/24 MATCH 10.82.105.5
```

Complex examples

The following examples are more complex and use a NetBackup server (`bob`), that uses the following network interfaces:

```
eri0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 10.82.105.11 netmask fffff800 broadcast 10.82.111.255

eri0:1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 10.82.105.8 netmask fffff800 broadcast 10.255.255.255

eri1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
    inet 10.82.10.10 netmask fffff800 broadcast 10.82.15.255
```

Normal outbound connectivity to the following four hosts (billcat, muzzy, beetle, lilo), uses the first interface. Internal connections use the destination interface as the source interface.

```
$ bptestbpcd -host billcat
10.82.105.11:54129 -> 10.82.105.15:13724

$ bptestbpcd -host muzzy
10.82.105.11:54152 -> 10.82.105.14:13724

$ bptestbpcd -host beetle
10.82.105.11:54135 -> 10.82.104.249:13724

$ bptestbpcd -host lilo
10.82.105.11:54139 -> 10.82.56.79:1556
$ bptestbpcd -host 10.82.105.11
10.82.105.11:54144 -> 10.82.105.11:1556
$ bptestbpcd -host 10.82.105.8
10.82.105.8:52148 -> 10.82.105.8:1556
```

Example 1

Using a local interface as the `target` for `MATCH` entries has no affect. In this example, the source interface is unaffected by the local `MATCH` entry.

```
PREFERRED_NETWORK = 10.82.105.8/32 MATCH

$ bptestbpcd -host billcat

10.82.105.11:54202 -> 10.82.105.15:13724

$ bptestbpcd -host muzzy
10.82.105.11:54206 -> 10.82.105.14:13724

$ bptestbpcd -host beetle
10.82.105.11:54300 -> 10.82.104.249:13724
```

```
$ bptestbpcd -host lilo
10.82.105.11:54302 -> 10.82.56.79:1556
$ bptestbpcd -host 10.82.105.11
10.82.105.11:54306 -> 10.82.105.11:1556
$ bptestbpcd -host 10.82.105.8
10.82.105.8:54309 -> 10.82.105.8:1556
```

Example 2

Similar to [Example 1](#), using a local interface as the `target` for `ONLY` entries has no affect on source binding. It does, however, prevent connections to destination addresses (in the absence of other directives that more closely `MATCH` the destinations). Connections internal to the host are not affected.

```
PREFERRED_NETWORK = 10.82.105.8/32 ONLY

$ bptestbpcd -host billcat
<16> bptestbpcd main: ConnectToBPCD(billcat) failed:
25 cannot connect on socket

$ bptestbpcd -host muzzy
<16> bptestbpcd main: ConnectToBPCD(muzzy) failed:
25 cannot connect on socket

$ bptestbpcd -host beetle
<16> bptestbpcd main: ConnectToBPCD(beetle) failed:
25 cannot connect on socket

$ bptestbpcd -host lilo
<16> bptestbpcd main: ConnectToBPCD(lilo) failed:
25 cannot connect on socket

$ bptestbpcd -host 10.82.105.11
10.82.105.11:54306 -> 10.82.105.11:1556

$ bptestbpcd -host 10.82.105.8
10.82.105.8:54309 -> 10.82.105.8:1556
```

Example 3

Using `MATCH` entries, the outbound connections to a specific host or network can be preferred over the defaults. In this example, connections to a specific host and a separate network are requested to use the second outbound network interface.


```
PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.8
PREFERRED_NETWORK = 10.82.104.0/24 MATCH 10.82.105.8

$ bptestbpcd -host billcat (Preferred by the first entry)
10.82.105.8:54192 -> 10.82.105.15:13724

$ bptestbpcd -host muzzy (Implicitly permitted using defaults)
10.82.105.11:54196 -> 10.82.105.14:13724

$ bptestbpcd -host beetle (Preferred by the second entry)
10.82.105.8:54200 -> 10.82.104.249:13724

$ bptestbpcd -host lilo (Implicitly permitted using defaults)
10.82.105.11:54202 -> 10.82.56.79:1556
```

Example 4

Adding an **ONLY** entry prevents connections to any other hosts that are not on the specified network, or matched by prior entries.

```
PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.8
PREFERRED_NETWORK = 10.82.104.0/24 MATCH 10.82.105.8
PREFERRED_NETWORK = 10.82.56.0/24 ONLY

$ bptestbpcd -host billcat (Preferred by first entry)
10.82.105.8:54209 -> 10.82.105.15:13724

<16> bptestbpcd -host 10.82.105.14 (Does not match 1 or 2, excluded by 3)
<16> bptestbpcd main: ConnectToBPCD(muzzy) failed: 25 cannot connect
on socket

$ bptestbpcd -host beetle (Preferred by second entry)
10.82.105.8:54214 -> 10.82.104.249:13724 (Required by third entry)

10.82.105.11:54216 -> 10.82.56.79:1556
```

Example 5

Changing the **ONLY** to **PROHIBITED** explicitly excludes connections with those destination hosts and implicitly allows connections to unspecified hosts. The **PROHIBITED** network is non-local and does not affect source binding.

```
PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.8
PREFERRED_NETWORK = 10.82.104.249/32 MATCH 10.82.105.8
PREFERRED_NETWORK = 10.82.56.0/24 PROHIBITED
```

```
$ bptestbpcd -host billcat (Preferred by the first entry)
10.82.105.8:54224 -> 10.82.105.15:13724

$ bptestbpcd -host muzzy (Implicitly permitted)
10.82.105.11:54228 -> 10.82.105.14:13724

$ bptestbpcd -host beetle (Preferred by the second entry)
10.82.105.8:54232 -> 10.82.104.249:13724

$ bptestbpcd -host 10.82.56.79 (Does not match 1 or 2, prohibited by 3)
<16> bptestbpcd main: ConnectToBPCD(lilo) failed: 25 cannot connect
on socket
```

Example 6

Conversely, moving the `ONLY` to the top of the list does not prevent the `MATCH` entries from being evaluated because the `ONLY` is for a less restrictive IP range than the `MATCH` entries. The latter are evaluated first for those hosts.

```
PREFERRED_NETWORK = 10.82.104.0/24 ONLY
PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.11
PREFERRED_NETWORK = 10.82.104.249/32 MATCH 10.82.105.8

$ bptestbpcd -host billcat (Preferred by the second entry)
10.82.105.11:54392 -> 10.82.105.15:13724

$ bptestbpcd -host 10.82.105.14 (Does not match 2 or 3, excluded by 1)
<16> bptestbpcd main: ConnectToBPCD(muzzy) failed: 25 cannot connect
on socket

$ bptestbpcd -host beetle (Preferred by 3 before required by 1)
10.82.105.8:54396 -> 10.82.104.249:13724

$ bptestbpcd -host 10.82.56.79 (Does not match 2 or 3, excluded by 1)
<16> bptestbpcd main: ConnectToBPCD(lilo) failed: 25 cannot connect
on socket
```

Example 7

The subnet on this `ONLY` entry matches both `billcat` and `muzzy`, but does not affect the outbound interface confirming that `ONLY` is used for destination address filtering and not source address filtering. Otherwise, all connections would fail because both local interfaces, `10.82.105.11` and `10.82.105.8`, are not in that subnet.

```
PREFERRED_NETWORK = 10.82.105.14/31 ONLY
PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.8

$ bptestbpcd -host billcat (Preferred by second entry)
10.82.105.8:54209 -> 10.82.105.15:13724

$ bptestbpcd -host muzzy (Preferred by first entry)
10.82.105.11:45662 -> 10.82.105.14:13724

$ bptestbpcd -host 10.82.104.249 (Excluded by first entry)
<16> bptestbpcd main: ConnectToBPCD(beetle) failed: 25 cannot connect
on socket
```

Example 8

Here, all three remote hosts are reachable, but notice that the source interface is the one remaining after 10.82.105.11 is `PROHIBITED`. This includes the apparent target `MATCH` for `billcat`, which actually failed to match because the `source` was previously `PROHIBITED`. Notice that internal connections are not affected by `PROHIBITED`.

```
PREFERRED_NETWORK = 10.82.105.11/32 PROHIBITED
PREFERRED_NETWORK = 10.82.105.15/32 MATCH 10.82.105.11

$ bptestbpcd -host billcat (Matched second, but first prohibited that source)
10.82.105.8:54202 -> 10.82.105.15:13724

$ bptestbpcd -host muzzy (Implicit match and pruned source)
10.82.105.8:54206 -> 10.82.105.14:13724

$ bptestbpcd -host beetle (Implicit match and pruned source)
10.82.105.8:54300 -> 10.82.104.249:13724

$ bptestbpcd -host 10.82.105.11 (Not affected by first entry)
10.82.105.11:54306 -> 10.82.105.11:1556
$ bptestbpcd -host 10.82.105.8
10.82.105.8:54309 -> 10.82.105.8:1556
```

Example 9

This example demonstrates two nuances of source binding evaluation that result in the use of `ANY` interface instead of the non-prohibited interfaces. The second entry removes the 10.82.10.10 local interface from the source binding list before the third entry is processed making that `source` unavailable. The `source` on the

first entry causes the shortened list created by the second entry to be ignored during all evaluations.

```
PREFERRED_NETWORK = 10.82.104.249 MATCH 10.82.105.0/24
PREFERRED_NETWORK = 10.82.10.10 PROHIBITED
PREFERRED_NETWORK = 10.82.56.0/24 MATCH 10.82.10.10
```

FL: billcat -> 10.82.105.15 ... SRC: ANY (First source implicitly negates second target)

FL: muzzy -> 10.82.105.14 ... SRC: ANY (First source implicitly negates second target)

FL: beetle -> 10.82.104.249 ... SRC: 10.82.105.11 (Matched first, used first in range)

FL: lilo -> 10.82.56.79 ... SRC: ANY (Second target explicitly negates third source)

In [Example 8](#), the `source` on the first entry matches two local interfaces. The 10.82.105.11 interface was chosen over 10.82.105.8 as the source when connecting to `beetle` because that interface was returned first by the operating system as shown in the `bplocaladdrs` output for this example. (See [the section called “Using bplocaladdrs to troubleshoot”](#) on page 290.)

Example 10

This example shows how the binding list is shortened by prohibiting a local interface. When `ANY` was the default source binding list, the outbound interface for these destinations was 10.82.105.11. (See [the section called “Example 1”](#) on page 291.) Prohibiting a different local interface causes NetBackup to provide a shortened list and the operating system selected 10.82.10.10 as the source IP. Because this operating system uses the strong host model, that interface is not valid for these destination IPs and the connection attempts fail.

```
PREFERRED_NETWORK = 10.82.105.8 PROHIBITED

FL:  billcat -> 10.82.105.15 ... SRC: 10.82.10.10,10.82.105.11
FL:    lilo -> 10.82.56.79   ... SRC: 10.82.10.10,10.82.105.11

$ bptestbpcd -host billcat
<16> bptestbpcd main: ConnectToBPCD(billcat) failed:
25 cannot connect on socket
$ bptestbpcd -host lilo
<16> bptestbpcd main: ConnectToBPCD(lilo) failed:
25 cannot connect on socket
```

If the operating system is changed to the weak host model, the TCP SYN for each connection is transmitted out the default interface (10.10.82.105.11) onto the 10.82.104.0 network, but with a source IP of 10.82.10.10. If there is a network route from the 10.82.104.0 network to the destination hosts, then the SYN will reach the destinations. But the reply is only successful if there is an asymmetrical route back to the 10.82.8.0 network from the destination host. Notice the spoofed source IP in the successful connection which does not reflect the network onto which the TCP SYN packet was actually sent.

```
$ bptestbpcd -host billcat
<16> bptestbpcd main: ConnectToBPCD(billcat) failed:
25 cannot connect on socket
$ bptestbpcd -host lilo
10.82.10.10:52842 -> 10.82.56.79:1556
```

Compatibility

Any legacy Required Interface or Required Network configuration is automatically converted to a Preferred Network representation internally.

Consider primary server `bob`, as described in a previous topic. (See [“PREFERRED_NETWORK examples”](#) on page 289.)

```
REQUIRED_INTERFACE = bob
```

This entry is equivalent to the following entry for IPv4:

```
PREFERRED_NETWORK = 0/0 MATCH 10.82.105.11
```

If IPv6 is enabled, using `IP_ADDRESS_FAMILY = AF_UNSPEC`, the following is equivalent for IPv6:

```
PREFERRED_NETWORK = 0/0 MATCH fc44:53f9:cb30:201:250:56ff:febc:e85f
```

Both bind the specified source interface for all outbound connections because 0/0 matches all destinations. But notice the length of subnet (/0): any other directive with a source binding and a longer target subnet will supersede these entries. Similarly, because both the IPv4 and IPv6 examples have the same subnet length, only the first of these two would be honored if both were configured.

Similarly, if a required network was configured:

```
REQUIRED_NETWORK = 10.82.105/21
```

It translates to the following:

```
PREFERRED_NETWORK = 10.82.105/21 ONLY
```

Which restricts destination addresses to the specified network without affecting source interface selection.

Notes:

- In the event that both `REQUIRED_INTERFACE` and `PREFERRED_NETWORK` are specified and if they conflict, `REQUIRED_INTERFACE` overrides.
- Unlike `REQUIRED_INTERFACE`, `PREFERRED_NETWORK` does not change the `requesting_client` or `destination_client` fields in user-directed requests to `bprd` for image list or restore.

RANDOM_PORTS option for NetBackup servers and clients

This option specifies whether NetBackup chooses port numbers randomly or sequentially when it requires one for communication with NetBackup on other computers.

Table 2-177 RANDOM_PORTS information

Usage	Description
Where to use	On NetBackup primary servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. ■ If <code>RANDOM_PORTS = YES</code> (default), NetBackup chooses port numbers randomly from those that are free in the allowed range. For example, if the range is from 1024 through 5000, it chooses randomly from the numbers in this range. ■ If <code>RANDOM_PORTS = NO</code>, NetBackup chooses numbers sequentially, starting with the highest number available in the allowed range. For example, if the range is from 1024 through 5000, NetBackup chooses 5000 (if the number is available). If 5000 is not available, port 4999 is chosen. By default, this option is not present in the configuration file and NetBackup uses the random method for selecting port numbers.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the host > Port ranges > Use random port assignments. See "Port ranges properties" on page 127.

RE_READ_INTERVAL option for NetBackup servers

The `RE_READ_INTERVAL` option determines how often NetBackup checks disk storage units for available capacity.

Table 2-178 `RE_READ_INTERVAL` information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RE_READ_INTERVAL = <i>seconds</i></pre> The default is 300 seconds (5 minutes). This entry should appear only once in the configuration file.
Example	The reread interval is changed to 15 minutes in the following example: <pre>RE_READ_INTERVAL = 900</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > General server > Check the capacity of disk storage units every. See “General server properties” on page 106.

REQUIRED_NETWORK option for NetBackup servers

The `REQUIRED_NETWORK` option specifies the required route for backup traffic in an environment where the network traffic is segregated.

For example, an environment can contain a production network at `145.21.14.0` and a backup network at `192.132.28.0`.

Table 2-179 `REQUIRED_NETWORK` information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-179 REQUIRED_NETWORK information (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>REQUIRED_NETWORK = IP_address</pre> This entry should appear only once in the configuration file. Note: If the variable is set and the network is not available, all connections fail and no backups are performed.
Example	The required network is set to <code>192.132.28.0</code> in the following example: <pre>REQUIRED_NETWORK = 192.132.28.0</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Preferred network > Only. See “Preferred network properties” on page 129.

RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH option for NetBackup servers

This directory path specifies where the media server processes write the job deferral caches. The deferral cache files are written to `/usr/opensv/tmp` on Linux, and `install_path\NetBackup\temp` on Windows by default. For the Flex Appliance, the deferral cache file path is set to `/mnt/nbstage/usr/opensv/tmp`.

NetBackup backups generate metadata relative to the amount of data that is backed up. If a media server has many active backups running in parallel, NetBackup can write a significant amount of data to the deferral cache files.

Be aware of the relationship between the `RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH` and `MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB` values. Use `RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH` to specify where to write the cache information. Use `MINIMUM_DEFERRAL_CACHE_FREE_SPACE_MB` to specify how much free disk space to maintain.

Table 2-180

Usage	Description
Where to use	On NetBackup media servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH = <i>path</i></pre> This option should appear only once in the configuration file.
Example	On the media server, add the following entry to specify where the media server processes write the job deferral caches: <pre>RESILIENT_BACKUP_JOB_DEFERRAL_CACHE_FILE_PATH = /var/cache</pre>

RESILIENT_BACKUP_JOB_RESTART_TIMEOUT option for NetBackup servers

This time-out value specifies how long the media server processes wait to reestablish the job after first losing communications with the primary server. If the primary server does not reestablish communications to media server processes within this time period, the backup fails. The value is specified in minutes. The default value is 60 minutes.

Table 2-181

Usage	Description
Where to use	On NetBackup media servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RESILIENT_BACKUP_JOB_RESTART_TIMEOUT = <i>number</i></pre> This option should appear only once in the configuration file.

Table 2-181 (continued)

Usage	Description
Example	On the media server, add the following entry to specify how long the media server processes should wait to reestablish the job: <pre>RESILIENT_BACKUP_JOB_RESTART_TIMEOUT = 30</pre>

RESILIENT_NETWORK option for NetBackup primary servers and clients

The `RESILIENT_NETWORK` option specifies the computers that should use a resilient connection for backups and restores.

Table 2-182 `RESILIENT_NETWORK` information

Usage	Description
Where to use	On NetBackup primary servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use any of the following formats: <pre>RESILIENT_NETWORK = hostname ON OFF</pre> <pre>RESILIENT_NETWORK = ip_address ON OFF</pre> <pre>RESILIENT_NETWORK = network address/network_mask ON OFF</pre> You can mix IPv4 addresses and ranges with IPv6 addresses and subnets. By default, <code>RESILIENT_NETWORK</code> is not present in the configuration file.
Examples	The following are examples of valid forms for this entry: <pre>RESILIENT_NETWORK = client.veritas.org ON</pre> <pre>RESILIENT_NETWORK = 192.0.2.0 ON</pre> <pre>RESILIENT_NETWORK = 192.0.2.0/26 OFF</pre> <pre>RESILIENT_NETWORK = 2001:db8:0:0:0:0:0:0 ON</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server or client > Resilient network. See “Specify resilient connections for clients” on page 149.

Note: The order is significant for the items in the list of resilient networks. If a client is in the list more than once, the first match determines its resilient connection status. For example, suppose you add a client and specify the client IP address and specify **On** for **Resiliency**. Suppose also that you add a range of IP addresses as **Off**, and the client IP address is within that range. If the client IP address appears before the address range, the client connection is resilient. Conversely, if the IP range appears first, the client connection is not resilient.

See “[NBRNTD_IDLE_TIMEOUT option for NetBackup servers](#)” on page 264.

RESILIENT_RECONNECT_TIMEOUT

This value is the maximum time the primary server should wait for `nbjm` to reestablish connections with an active backup job after it encounters a network error. The time is specified in seconds. The default value is 600, or 10 minutes. The minimum value is 30 and the maximum value is 3600, or 1 hour.

Be aware of the relationship between `RESILIENT_RECONNECT_TIMEOUT` and `RESILIENT_RETRY_INTERVAL`. For example, if `RESILIENT_RECONNECT_TIMEOUT` is 10 minutes and `RESILIENT_RETRY_INTERVAL` is 2 minutes, the primary server attempts to reconnect to a backup job 5 times. If the primary server is unable to reconnect, the primary server ends the backup job.

Table 2-183

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RESILIENT_RECONNECT_TIMEOUT = number</pre> This option should appear only once in the configuration file.
Example	On the primary server, add the following entry to specify the maximum time <code>nbjm</code> should wait to reconnect to a backup job after it encounters a network error: <pre>RESILIENT_RECONNECT_TIMEOUT = 1800</pre>

RESILIENT_RETRY_INTERVAL

This value determines how frequently the primary server attempts to reestablish communications with an active backup job after it encounters a network error. Once the primary server encounters a network error, it waits the specified amount of time before it attempts to reestablish communications. The time is specified in seconds. The default value is 120, or 2 minutes. The minimum value is 1 and the maximum value is 3600, or one hour.

Be aware of the relationship between `RESILIENT_RECONNECT_TIMEOUT` and `RESILIENT_RETRY_INTERVAL`. For example, if `RESILIENT_RECONNECT_TIMEOUT` is 10 minutes and `RESILIENT_RETRY_INTERVAL` is 2 minutes, the primary server attempts to reconnect to a backup job 5 times. If the primary server is unable to reconnect, the primary server ends the backup job.

Table 2-184

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RESILIENT_RETRY_INTERVAL = number</pre> This option should appear only once in the configuration file.
Example	On the primary server, add the following entry to specify the minimum time to retry a request that is resiliency protected: <pre>RESILIENT_RETRY_INTERVAL = 360</pre>

RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers

The `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL` option specifies that NetBackup should perform normal duplication if an OpenStorage optimized duplication fails.

Table 2-185 `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL` information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-185 RESUME_ORIG_DUP_ON_OPT_DUP_FAIL information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE FALSE</pre> By default, this entry is not present and NetBackup does not perform normal duplication when an optimized duplication fails. This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

REVERSE_NAME_LOOKUP option for NetBackup servers and clients

This option lets administrators allow, restrict, or prohibit reverse host name lookup.

Table 2-186 REVERSE_NAME_LOOKUP information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>REVERSE_NAME_LOOKUP = ALLOWED RESTRICTED PROHIBITED</pre> This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the host > Network settings > Reverse host name lookup. See “Network settings properties” on page 124.

SECURE_PROXY_CIPHER_LIST option for NetBackup servers and clients

The `SECURE_PROXY_CIPHER_LIST` option configures the ciphers that NetBackup uses for OpenSSL to encrypt communication through the `vnetd` network daemon. The `SECURE_PROXY_CIPHER_LIST` option is a colon-separated list of permitted OpenSSL cipher strings. For the permitted cipher strings, requirements, and limitations, see the OpenSSL cipher documentation.

You can use this option to change the ciphers that NetBackup uses. If you configure this option, NetBackup writes a message about your configured cipher strings to the `vnetd nbpxyhelper VxUL` logs. The following is an example:

```
"Using user configured cipher list: cipher_string:cipher_string:...
```

Warning: Be careful when you configure the `SECURE_PROXY_CIPHER_LIST` option. Permitted OpenSSL lower-level primitives may overlap with the ciphers that provide no authentication or no encryption. Hosts that do not have a cipher in common in their cipher lists cannot communicate with each other.

Table 2-187 `SECURE_PROXY_CIPHER_LIST` information

Usage	Description
Where to use	On NetBackup primary servers, media servers, or clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SECURE_PROXY_CIPHER_LIST = cipher_string:cipher_string:cipher_string:...</pre> Replace <code>cipher_string</code> with a permitted OpenSSL cipher string. By default, the <code>SECURE_PROXY_CIPHER_LIST</code> option is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

SERVER option for NetBackup servers

The first `SERVER` option in the `bp.conf` file must point to the primary server where the `bp.conf` option resides. During installation, `SERVER` is automatically set to the name of the system where the NetBackup primary server software is installed.

Table 2-188 SERVER information

Usage	Description
Where to use	On NetBackup primary servers and media servers. Note: For a Fibre Transport (FT) media server that has multiple network interfaces for VLANs, ensure that the FT server's primary host name appears before any other interface names for that FT media server host. For information about these commands, see the NetBackup SAN Client and Fibre Transport Guide.
How to use	An entry for the <code>SERVER</code> option must be present in the configuration file on all NetBackup servers and clients. It is the only required NetBackup option. This option is not used in <code>\$HOME/bp.conf</code> files on a client. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Note: This topic discusses the <code>bp.conf</code> entries on the server. However, note that every <code>SERVER</code> option in a client <code>bp.conf</code> file must be a NetBackup primary or media server. That is, each system that is listed as a <code>SERVER</code> must have either NetBackup primary or media server software installed. The client service on some clients cannot be started if the client name is incorrectly listed as a server. If you configure NetBackup media servers for a primary server, the <code>bp.conf</code> file on the primary server must have a <code>SERVER</code> entry or <code>MEDIA_SERVER</code> entry for each. As previously mentioned, the first <code>SERVER</code> entry in the list designates the primary server itself. The <code>SERVER</code> entry or the <code>MEDIA_SERVER</code> entries should be added after the first, self-referencing option. A NetBackup primary server can be backed up as a NetBackup client by the servers that belong to another cluster. In that case the <code>bp.conf</code> file on the primary server should have <code>SERVER</code> entries for those servers as well.

Table 2-188 SERVER information (*continued*)

Usage	Description
Example	The following is an example entry on a primary server: <pre>SERVER = Primary_server (this primary server itself) SERVER = NB_server (primary server of another cluster) SERVER = Media_server_#1 MEDIA_SERVER = Media_server_#2 . . .</pre> The first <code>SERVER</code> entry on all the media servers must point to the primary server for those media servers. A media server can have only one primary server. However, a media server can be backed up as a NetBackup client by the servers that belong to another cluster. In this case, the configuration file on the media server must have <code>SERVER</code> entries for those servers as well. The following is an example entry on a media server: <pre>SERVER = Primary_server (for this media server) SERVER = NB_server (primary server of another cluster) SERVER = Media_server_#1 MEDIA_SERVER = Media_server_#2 . . .</pre> The <code>SERVER</code> entries must be the same on all servers in a primary and a media server cluster. If a <code>SERVER</code> entry is added or modified in the <code>bp.conf</code> file on the primary server, stop and restart <code>bprd</code> and <code>bpdcm</code> so that NetBackup recognizes the change. (The NetBackup request daemon and NetBackup Database Manager.) Note: If the first <code>SERVER</code> entry (the primary server) is modified on a media server, the NetBackup database also needs to be updated. Run <code>nbenmcmd -updatehost</code> to change the primary server for a media server.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Servers. See “Servers properties” on page 159. See “<code>FAILOVER_RESTORE_MEDIA_SERVERS</code> option for NetBackup servers” on page 237.

SERVER_CONNECT_TIMEOUT option for NetBackup servers

`SERVER_CONNECT_TIMEOUT` specifies the number of seconds that the primary server waits before it times out when it connects to a media server.

Table 2-189 `SERVER_CONNECT_TIMEOUT` information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SERVER_CONNECT_TIMEOUT = seconds</pre> The default timeout period is 30 seconds. This entry should appear only once in the configuration file.
Example	The example permits a timeout of 60 seconds: <pre>SERVER_CONNECT_TIMEOUT = 60</pre>

SERVER_PORT_WINDOW option for NetBackup servers

The `SERVER_PORT_WINDOW` option specifies the range of non-reserved ports on which NetBackup processes on this computer accept connections from NetBackup on other computers when the inbound connection is not to a well known port. This primarily applies to `bpcd` call-back when `vnetd` is disabled in the connect options for the remote NetBackup server or client and that host is configured for non-reserved ports. This also applies to NDMP call-back to the media server during remote NDMP backups.

Table 2-190 `SERVER_PORT_WINDOW` information

Usage	Description
Where to use	On NetBackup primary servers or media servers.

Table 2-190 SERVER_PORT_WINDOW information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SERVER_PORT_WINDOW = start_port end_port</pre> The default range is from 1024 through 5000. This entry should appear only once in the configuration file.
Example	The following example permits ports from 4900 through 5000: <pre>SERVER_PORT_WINDOW = 4900 5000</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Port ranges > Server port window. See "Port ranges properties" on page 127.

SERVER_RESERVED_PORT_WINDOW option for NetBackup servers and clients

This option specifies the range of local reserved ports on which this computer accepts connections from NetBackup on other computers when the inbound connection is not to a well known port. This primarily applies to `bpcd` call-back when `vnetd` is disabled in the connect options for the remote NetBackup server or client.

The `SERVER_RESERVED_PORT_WINDOW` option applies when a server connects to a client that is configured to accept only reserved ports. This option is generally not useful on clients.

Table 2-191 SERVER_RESERVED_PORT_WINDOW information

Usage	Description
Where to use	On NetBackup primary servers or media servers.

Table 2-191 SERVER_RESERVED_PORT_WINDOW information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SERVER_RESERVED_PORT_WINDOW = start_port end_port</pre> The default range is from 512 through 1023. This entry should appear only once in the configuration file.
Example	The following example permits ports from 900 through 1023: <pre>SERVER_RESERVED_PORT_WINDOW = 900 1023</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Port ranges > Server reserved port window. See “Port ranges properties” on page 127.

SKIP_RESTORE_TO_SYMLINK_DIR option for NetBackup servers

The `SKIP_RESTORE_TO_SYMLINK_DIR` option forces NetBackup to check all directories on a UNIX client into which files are restored. If the file to be restored is under a symbolically linked directory, NetBackup does not restore the file.

Table 2-192 SKIP_RESTORE_TO_SYMLINK_DIR information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-192 SKIP_RESTORE_TO_SYMLINK_DIR information (*continued*)

Usage	Description
How to use	This option should appear only once in the configuration file. For example, if the UNIX client requests a restore for <code>/home/user/.cshrc</code> and <code>/home/user</code> is a symbolic link, NetBackup does not restore <code>.cshrc</code>. The addition of <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> helps minimize potential security and data loss problems if the restore is performed with root permissions. Without <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> in the <code>bp.conf</code> file, NetBackup follows any symbolically linked directories and restores files to that location. Note: Restore job performance is reduced by using this option. <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> and <code>UNLINK_ON_OVERWRITE</code> do not affect each other if both are specified, with one exception: When the following options are enabled: ■ Overwrite existing files option ■ <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> ■ <code>UNLINK_ON_OVERWRITE</code> Then, when a restore job comes across a symbolic link, the link is unlinked before the job checks, and the files and directory are restored. For example, <code>/home/user/</code> is backed up as a directory and, when restored, it is a symbolic link to a directory. These settings have the following outcomes: ■ With <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> set (and Overwrite existing files indicated), no files are restored into the directory the symbolic link points to, and the symbolic link remains. ■ With both <code>UNLINK_ON_OVERWRITE</code> and <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> (and Overwrite existing files indicated), the symbolic link directory is unlinked, the original directory is restored, and all files within the directory are also restored. ■ With neither option set (and Overwrite existing files indicated), NetBackup follows the symbolic link and restore all files into the directory to which the symbolic link points.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the primary server > Port ranges > Server reserved port window. See “Port ranges properties” on page 127.

SYSLOG_AUDIT_CATEGORIES for NetBackup primary server

Use the `SYSLOG_AUDIT_CATEGORIES` option to send the NetBackup audit events to system logs. You can view NetBackup audit events in the system logs. For example,

on a Windows system, use Windows Event Manager to view NetBackup audit events.

Table 2-193 SYSLOG_AUDIT_CATEGORIES information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SYSLOG_AUDIT_CATEGORIES = audit_category1, audit_category2</pre> For example, to send events of the POLICY and JOB audit categories to the system logs, use the following format: <pre>SYSLOG_AUDIT_CATEGORIES = POLICY,JOB</pre> To send events of all audit categories to the system logs, use the following format: <pre>SYSLOG_AUDIT_CATEGORIES = ALL</pre>
Equivalent NetBackup web UI property	Security > Security events > Security event settings> Send the audit events to the system logs.

SYSLOG_AUDIT_USE_OCSF_FORMAT for NetBackup primary server

By default, NetBackup sends the audit events to system logs in native format.

Use the `SYSLOG_AUDIT_USE_OCSF_FORMAT` option to send the NetBackup audit events to system logs in the Open Cybersecurity Schema Framework (OCSF) format.

You can view NetBackup audit events in the system logs. For example, on a Windows system, use Windows Event Manager to view NetBackup audit events.

Table 2-194 SYSLOG_AUDIT_USE_OCSF_FORMAT

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-194 SYSLOG_AUDIT_USE_OCSF_FORMAT (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SYSLOG_AUDIT_USE_OCSF_FORMAT = TRUE</pre>
Equivalent NetBackup web UI property	No equivalent property on the web UI

TELEMETRY_UPLOAD option for NetBackup servers

`TELEMETRY_UPLOAD` allows NetBackup to collect data about how the product is used in a NetBackup environment. The information becomes part of a continuous quality improvement program (NetBackup Product Improvement Program) that helps the NetBackup development and support teams understand how customers configure, deploy, and use the NetBackup product. The data is used for only product development and problem analysis purposes in the company.

The telemetry collection agent runs on every NetBackup server. The telemetry collection agent does not run on NetBackup clients.

Note: If Cohesity Alta manages the server, this value is ignored.

Table 2-195 TELEMETRY_UPLOAD information

Usage	Description
Where to use	On NetBackup primary and media servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>TELEMETRY_UPLOAD = YES NO</pre> The default is <code>YES</code>. This entry should appear only once in the configuration file.

THROTTLE_BANDWIDTH option for NetBackup servers

This option specifies a limit for the network bandwidth or transfer rate that NetBackup clients use on a network. The actual limiting occurs on the client side of the backup connection. This option limits only backups. Restores are unaffected. The default is that the bandwidth is not limited.

While `LIMIT_BANDWIDTH` associates a bandwidth or transfer rate with all client IP addresses in a range between two IP addresses, `THROTTLE_BANDWIDTH` is more useful in an IPv6 environment. `THROTTLE_BANDWIDTH` associates a bandwidth setting with a subnet description.

For example, the following subnet will get 400kbs bandwidth:

```
2001:db8:cb30:120::/64 400
```

Table 2-196 THROTTLE_BANDWIDTH information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format in a single line: <pre>THROTTLE_BANDWIDTH = xxxx:xxxx:xxxx:xxxx: yyyy:yyyy:yyyy:yyyy::/nnn zzz</pre> Each <code>THROTTLE_BANDWIDTH</code> option specifies the bandwidth value and the IP address of the clients and networks to which it applies. The following variables are defined: ■ <code>xxxx.xxxx.xxxx.xxxx</code> is the subnet portion (64-bits) of the IPv6 address range. (For example, <code>2001:db8:1:110</code>.) ■ <code>yyyy.yyyy.yyyy.yyyy</code> is the host portion (64-bits) of the IPv6 address range. (For example, <code>0:0:0:8b72</code>.) ■ <code>nnn</code> is the number of mask bits that, when applied to the IPv6 address, identifies the range of addresses that are considered for throttling. The valid range is 0 to 128. Mask bits are applied left to right across the address range. ■ <code>zzz</code> is the bandwidth limitation in kilobytes per second. (For example, <code>200</code>.) A value of 0 disables throttling IPv6 addresses covered by this option.

Table 2-196 THROTTLE_BANDWIDTH information (continued)

Usage	Description
Bandwidth examples	The following are <code>LIMIT_BANDWIDTH</code> examples: ■ Configure a bandwidth limit of 500 kilobytes per second for all computers on the subnet 2001:db8:1:110 as follows: <code>LIMIT_BANDWIDTH = 2001:db8:1:110::/64 500</code> ■ Configure a bandwidth limit of 700 kilobytes per second for a particular client (2001:db8:1:110:0:0:0:8b72) as follows: <code>LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 700</code> ■ To disable bandwidth limiting for a client in a subnet that has a bandwidth limit, specify 0 for the kilobytes per second: <code>LIMIT_BANDWIDTH = 2001:db8:1:110::/64 500</code> <code>LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 0</code> In this case, no limiting occurs for the client with IPv6 address 2001:db8:1:110:0:0:0:8b72

IPv6 address rules for NetBackup clients

The IPv6 address ranges can specify individual clients or entire subnets.

- An IPv6 address can take the following forms:
 - a.b.c.d.
Where a, b, c, and d are hexadecimal integers in the range 0-ffff.
 - a
A 32-bit integer that represents the full IP address in network byte order. (The big endian, the most significant byte is first on the wire.)
- Enter IPv6 addresses as hexadecimal numbers.
- Neither the subnet nor the host part of an IPv6 address can be zero.
- Only ordinary IPv6 addresses are accepted.
- Do not create multiple entries that specify the same range of IPv6 addresses. If multiple entries are created, NetBackup uses the last one found. In the following example, NetBackup uses the second entry:

```
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 200
```

This rule also applies to multiple entries that specify an exact client address, as follows:


```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 200
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 100
```

- Do not specify IPv6 address ranges that overlap one another.
 Consider the following:

```
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
```

The ranges overlap, and bandwidth limiting results are unpredictable.

- Specify a range of addresses in one entry and an address for a specific client in other entries.

If a client is covered by an entry that specifies its exact IPv6 address and by another entry that specifies a range of IPv6 addresses, NetBackup uses the bandwidth value in the entry with the exact IP address.

The following sets the bandwidth for a range of IPv6 addresses:

```
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
```

The following sets the bandwidth for a specific address that is within the range:

```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 200
```

In this case, NetBackup uses the specific entry (bandwidth of 200) for the client whose address is 2001:db8:1:110:0:0:0:8b72. This capability can also be used to exclude specific clients from bandwidth limiting. The order of the range and specific address entries in the `bp.conf` file is not significant.

Rules for setting bandwidth values for NetBackup clients

Set bandwidths for individual clients to one of the following values:

- 0 (no bandwidth limiting), or
- Less than or equal to any value that is set for the IPv6 address range that contains the IP address for the client.

For example, the following is valid:

```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/64 500
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 300
```

If the bandwidth is set higher for a client than is set for the range, NetBackup ignores the individual setting. NetBackup uses the value for the range instead. In this case, the client receives a share of the bandwidth that is specified for the network.

If the bandwidth limit for a client is equal to or lower than the value for the range, the client uses the lower of the following settings:

- Its share of the network bandwidth value.
- Its individual bandwidth value.

The bandwidth value that NetBackup uses for a client is always at least one kilobyte per second.

TRUSTED_PRIMARY option for NetBackup servers

The `TRUSTED_PRIMARY` option lets administrators indicate a specific storage lifecycle policy in a target primary server domain to configure Auto Image Replication. The ability to replicate to a specific target domain SLP is supported between MSDP storage servers and PDDO storage servers.

Table 2-197 `TRUSTED_PRIMARY` information

Usage	Description
Where to use	On NetBackup primary servers or media servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Add <code>TRUSTED_PRIMARY</code> on the primary server in the source domain and the primary server in the target domain.
Example	On the source domain primary server, enter the name of the target primary server: <pre>TRUSTED_PRIMARY = Target_Server_Name</pre> On the target domain primary server, enter the name of the source primary server: <pre>TRUSTED_PRIMARY = Source_Server_Name</pre>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server > Servers > Trusted primary servers tab. Settings > Global security > Trusted primary servers See “Servers properties” on page 159.

ULINK_ON_OVERWRITE option for NetBackup servers

When a UNIX client indicates **Overwrite existing files** as a restore option, the `UNLINK_ON_OVERWRITE` option forces NetBackup to perform the following actions:

- Check for the existence of a file to be restored.
- Unlink the file if it exists.
- Restore the file.

The file can be any normal file, symbolic link, hard link, or empty directory.

The addition of `UNLINK_ON_OVERWRITE` helps minimize potential security and data loss problems from following existing symbolic links. It also guarantees that files are restored exactly as they were backed up.

Table 2-198 `UNLINK_ON_OVERWRITE` information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	This option should appear only once in the configuration file. Note: Restore job performance is reduced by using this option. If the <code>UNLINK_ON_OVERWRITE</code> option is not indicated in the <code>bp.conf</code> file but the Overwrite existing files option is specified, the behavior of NetBackup is different when it restores symbolic links. (Or, if the option is set to NO.) NetBackup unlinks existing files or empty directories when it restores symbolic links, hard links, or special files (CHR, BLK, and FIFO). However, NetBackup does not unlink when it restores normal files or directories, which can be problematic. NetBackup follows the symbolic link and creates or replaces the files that the symbolic link points to. Or, it replaces the directory that the symbolic link points to. <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> and <code>UNLINK_ON_OVERWRITE</code> do not affect each other if both are specified, with one exception: When the following options are enabled: ■ Overwrite existing files option ■ <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> ■ <code>UNLINK_ON_OVERWRITE</code> Then, when a restore job comes across a symbolic link, the link is unlinked before the job checks, and the files and directories are restored.

Table 2-198 UNLINK_ON_OVERWRITE information (*continued*)

Usage	Description
Example	For example, if the <code>/home/user/</code> path was backed up as a directory and, when restored, it is a symbolic link to a directory: ■ With <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> set (and Overwrite existing files indicated), no files are restored into the directory the symbolic link points to, and the symbolic link remains. ■ With both <code>UNLINK_ON_OVERWRITE</code> and <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> (and Overwrite existing files indicated), the symbolically linked directory is unlinked, the original directory is restored, and all files within the directory are also restored. ■ With neither set (and Overwrite existing files indicated), NetBackup follows the symbolic link and restore all files into the directory the symbolic link points to.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

USE_URANDOM for NetBackup servers and clients

In computing, entropy is the randomness collected by an operating system or application for use in cryptography or other uses that require random data.

Enable the `USE_URANDOM` option to specify `/dev/urandom` as the character device to provide cryptographically secure random output in your NetBackup environment.

Table 2-199 USE_URANDOM information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 2-199 USE_URANDOM information (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. The default value of the <code>USE_URANDOM</code> option is 0. When the <code>USE_URANDOM</code> option is set to default, the character device to be used is based on the value of the <code>NB_FIPS_MODE</code> option. If <code>NB_FIPS_MODE</code> is enabled, <code>dev/random</code> is used. If <code>NB_FIPS_MODE</code> is disabled, <code>dev/urandom</code> is used. See “NB_FIPS_MODE option for NetBackup servers and clients” on page 263. To enable the <code>USE_URANDOM</code> option, use the following format: <pre>USE_URANDOM = 1</pre> If <code>USE_URANDOM</code> is set to 2 (or is disabled), the <code>dev/random</code> character device is used to provide cryptographically secure random output.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

VERBOSE option for NetBackup servers and clients

Used for debugging purposes, the `VERBOSE` option controls the amount of information NetBackup includes in its legacy logs.

Table 2-200 VERBOSE information

Usage	Description
Where to use	On NetBackup primary servers or clients.

Table 2-200 VERBOSE information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>VERBOSE = [0 1 2 3 4 5]</pre> By default, this option is disabled. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the server or client > Logging > Global logging level. See “Logging properties” on page 113.

VIRTUALIZATION_CRL_CHECK for NetBackup servers and clients

The `VIRTUALIZATION_CRL_CHECK` option lets you specify the revocation check level for external certificates of the virtualization server. Based on the check, revocation status of the virtualization server certificate is validated against the certificate revocation list (CRL) during host communication.

By default, the `VIRTUALIZATION_CRL_CHECK` option is disabled. If you want to validate the revocation status of the virtualization server certificate against certificate revocation list (CRL), set the option to a different value.

You can choose to use the CRLs from the directory that is specified for the `ECA_CRL_PATH` configuration option or the CRL distribution point (CDP).

See [“ECA_CRL_PATH for NetBackup servers and clients”](#) on page 224.

Table 2-201 VIRTUALIZATION_CRL_CHECK information

Usage	Description
Where to use	On NetBackup primary server or all access hosts.

Table 2-201 **VIRTUALIZATION_CRL_CHECK** information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>VIRTUALIZATION_CRL_CHECK = CRL check</pre> You can specify one of the following: ■ DISABLE (or 0) - Revocation check is disabled. Revocation status of the certificate is not validated against the CRL during host communication. This is the default value. ■ LEAF (or 1) - Revocation status of the leaf certificate is validated against the CRL. ■ CHAIN (or 2) - Revocation status of all certificates from the certificate chain are validated against the CRL.
Equivalent NetBackup web UI property	No equivalent exists.

VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED for servers and clients

The `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` option enables the validation of virtualization server certificates using its root or intermediate certificate authority (CA) certificates.

Before you enable the option, review the steps from the 'Validating VMware virtualization server certificates in NetBackup' section in the [NetBackup for VMware Administrator's Guide](#).

By default, the `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` option is set to `UNDEFINED`.

The security certificate validation is enabled for Red Hat Virtualization and Nutanix AHV servers, but is disabled for VMware servers.

Note: In a scenario where an external CA can be configured for one virtualization server, but not for the other, two separate backup hosts must be used. The `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` option must be set to `YES` for the backup host where the external CA can be configured. The `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` must be set to `YES` for the backup host where the external CA can be configured. The option must be set to `NO` for the other backup host.

Table 2-202 `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` information

Usage	Description
Where to use	On NetBackup primary server or all access hosts.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format to enable certificate validation for the Red Hat Virtualization, VMware, or Nutanix AHV servers: <pre>VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED = YES</pre>
Equivalent NetBackup web UI property	No equivalent exists.

VIRTUALIZATION_HOSTS_CONNECT_TIMEOUT for servers and clients

The `VIRTUALIZATION_HOSTS_CONNECT_TIMEOUT` option lets you specify the duration (in seconds) after which the connection between NetBackup and vCloud Director server ends.

Table 2-203

Usage	Description
Where to use	On NetBackup primary server or all access hosts.

Table 2-203 (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. By default, the <code>VIRTUALIZATION_HOSTS_CONNECT_TIMEOUT</code> option is set to 60 seconds. Use the following format to specify the time-out value: <pre>VIRTUALIZATION_HOSTS_CONNECT_TIMEOUT = Time-out value in seconds</pre>
Equivalent NetBackup web UI property	No equivalent exists.

VMWARE_AUTODISCOVERY_INTERVAL option for NetBackup servers

This option controls how often NetBackup scans the vCenter servers to discover virtual machines to display in the NetBackup web UI.

NetBackup attempts autodiscovery first with the same host for which the last discovery attempt was successful. If autodiscovery fails with that host, NetBackup tries again with other hosts in the following order:

- The NetBackup primary server
- The access host, client, or proxy server
- The media server

Table 2-204 VMWARE_AUTODISCOVERY_INTERVAL information

Usage	Description
Where to use	On NetBackup primary servers.

Table 2-204 VMWARE_AUTODISCOVERY_INTERVAL information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Note: These commands require administrator privilege on the NetBackup primary server. For assistance, contact the NetBackup administrator. You can also use the NetBackup configuration APIs to view, add, or change this option. Refer to the NetBackup API documentation on SORT for more information. The default is 8 hours. The minimum is 5 minutes, the maximum 1 year. If set to zero, autodiscovery is disabled for all the VMware servers. Use the following format: <pre>VMWARE_AUTODISCOVERY_INTERVAL = number of seconds</pre> For example: <pre>VMWARE_AUTODISCOVERY_INTERVAL = 100000</pre> This entry should appear only once in the configuration file. Note: After changing this option, stop and restart the NetBackup services. For VM discovery, the <code>Netbackup Discovery Framework</code> service must be running.
Equivalent NetBackup web UI property	Workloads > VMware > VMware settings > Autodiscovery

VM_SNAPSHOT_QUIESCE_STATUS option for NetBackup servers and clients

NetBackup offers the option to perform a “quiescence failover” and then perform an unquiesced snapshot if the quiescence of a target VMware virtual machine fails. By default the snapshot job reports as status 0 (Success).

The `VM_SNAPSHOT_QUIESCE_STATUS` option provides the ability to change the status of the snapshot job from 0 to 1 (Partial success) when a quiescence of a target VMware virtual machine fails.

Table 2-205 VM_SNAPSHOT_QUIESCE_STATUS information

Usage	Description
Where to use	On the NetBackup media server and the backup host.
How to use	By default, this option is not present in the configuration file or registry. In that case the snapshot job has a status of 0 (Successful) for an unquiesced snapshot if a quiescence of a target VMware virtual machine fails. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Note: These commands require administrator privilege on the NetBackup primary server. For assistance, contact the NetBackup administrator. You can also use the NetBackup configuration APIs to view, add, or change this option. Refer to the NetBackup API documentation on SORT for more information. Use the following format: <pre>VM_SNAPSHOT_QUIESCE_STATUS = polycyname1, polycyname2, polycyname3</pre> For example: <pre>VM_SNAPSHOT_QUIESCE_STATUS = vmware1, vmware2</pre> This entry should appear only once in the configuration file or registry. Note: After changing this option, stop and restart the NetBackup services.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

VMWARE_TLS_MINIMUM_V1_2 for NetBackup servers and clients

The `VMWARE_TLS_MINIMUM_V1_2` option lets you specify the Transport Layer Security (TLS) version to be used for communication between NetBackup and VMware servers.

Table 2-206

Usage	Description
Where to use	On NetBackup primary server or all access hosts.

Table 2-206 (continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>VMWARE_TLS_MINIMUM_V1_2 = YES NO</pre> By default, the <code>VMWARE_TLS_MINIMUM_V1_2</code> option is set to <code>YES</code>. If the option is set to <code>YES</code>, TLS 1.2 version and the following cipher suites are used for communication with VMware servers: <pre>TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_RSA_WITH_AES_256_GCM_SHA384, TLS_DHE_DSS_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_RSA_WITH_AES_128_GCM_SHA256, TLS_DHE_DSS_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_RSA_WITH_AES_128_GCM_SHA256</pre> If the option is set to <code>NO</code>, the TLS 1.2 or earlier version with the default cipher suite is used for communication between NetBackup and VMware server. The cipher suite is used based on the TLS and cipher suite configuration that are set on the respective VMware server.
Equivalent NetBackup web UI property	No equivalent exists.

WEB_SERVER_TUNNEL_ENABLE option for NetBackup servers

In a demilitarized zone (DMZ), the client tries to communicate with the primary directly, and if the connection fails, tries to communicate using an HTTP tunnel on the media server. You can use the `WEB_SERVER_TUNNEL_ENABLE` option to disable the HTTP tunnel on a specific media server. You can use this option if the media server takes a backup of clients that are not in a DMZ.

For more information, refer to the **About the communication between a NetBackup client located in a demilitarized zone and a primary server through an HTTP tunnel** section in the *NetBackup Security and Encryption Guide*.

Table 2-207 WEB_SERVER_TUNNEL_ENABLE information

Usage	Description
Where to use	On NetBackup media servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>WEB_SERVER_TUNNEL_ENABLE = 1 0</pre> The option uses the following parameters: ■ 1 for using the HTTP tunnel. This value is considered as default. ■ 0 for disabling the communication using the HTTP tunnel. This value ensures that the HTTP tunnel process does not start. By default, the WEB_SERVER_TUNNEL_ENABLE option is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

Configuration options for NetBackup clients

The following topics are about configuration options for NetBackup clients. Nearly all of these options can also be set in the Host properties in the **NetBackup web UI**.

ACCEPT_REVERSE_CONNECTION for clients

The `ACCEPT_REVERSE_CONNECTION` option lets you start the subscriber service on a NAT client or a NAT server.

Once the option is enabled, you must restart the client services.

Table 2-208 ACCEPT_REVERSE_CONNECTION information

Usage	Description
Where to use	On NetBackup clients.

Table 2-208 ACCEPT_REVERSE_CONNECTION information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>ACCEPT_REVERSE_CONNECTION = TRUE</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

APP_PROXY_SERVER option for NetBackup clients

The `APP_PROXY_SERVER` entry specifies the name of the client as a backup host for allowed listing.

Table 2-209 APP_PROXY_SERVER information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>APP_PROXY_SERVER = clientname</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

BACKUP_BTRFS_SNAPSHOT option for NetBackup clients

This option indicates that the contents of the BTRFS file system (BTRFS) snapshots are to be backed up.

Table 2-210 BACKUP_BTRFS_SNAPSHOT information

Usage	Description
Where to use	On NetBackup clients.
How to use	By default, <code>BACKUP_BTRFS_SNAPSHOT</code> is not present in the configuration file. When the option is not enabled, only the BTRFS subvolumes are backed up, but not the BTRFS snapshots. When the option is present and enabled (1), the contents of the BTRFS snapshots on the client are backed up. Use the following format: <pre>BACKUP_BTRFS_SNAPSHOT = 1</pre> This entry should appear only once in the configuration file.
Example	The following format ensures that the contents of the BTRFS snapshots are backed up: <pre>BACKUP_BTRFS_SNAPSHOT = 1</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.
Additional information	The option is applicable for only BTRFS. By default, the contents of the BTRFS snapshots are not backed up. To ignore the snapshot paths, the <code>libbtrfsutil</code> package (version 4.17 onwards) must be present on the client computer. If the <code>libbtrfsutil</code> package is not present on the client computer, the read-only snapshot or subvolume backup is ignored.

BACKUP_FIFO_FILES option for NetBackup clients

This option indicates that the contents of a named pipe are to be backed up. A named pipe—also known as a FIFO—is a method of Inter-Process Communication that uses the file system interface to transfer data.

Table 2-211 BACKUP_FIFO_FILES information

Usage	Description
Where to use	On NetBackup clients.

Table 2-211 BACKUP_FIFO_FILES information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. By default, <code>BACKUP_FIFO_FILES</code> is not present in the configuration file. When the option is not enabled, only the metadata of the named pipe files is backed up, but not the data in the files. When the option is present and enabled (1), the contents of the named pipe files on the client are backed up. Use the following format: <pre>BACKUP_FIFO_FILES = 1</pre> This entry should appear only once in the configuration file.
Example	The following format ensures that the contents of the named pipe files on the client are backed up: <pre>BACKUP_FIFO_FILES = 1</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.
Additional information	For the contents of the files to be backed up, the write process must close the pipe. As long as data is passing through the pipe, the backup continues. If the third-party process writing to the named pipe file does not close the pipe, NetBackup fails the backup job with a non-zero status code (typically 13 or 41). Server resources are then allocated to other jobs. In the case of restore jobs, the job fails with status code 2800. In either case, the client processes continue waiting until they are terminated. To back up the script or executable that is responsible for writing to the named pipe, add the program file instead of the named pipe file to the Backup selections list. To avoid backing up named pipes, but still having the contents of other pipes read, add the named pipe to an exclude list or avoid the files entirely in the Backup selections list.

BPARCHIVE_POLICY option for NetBackup clients

The `BPARCHIVE_POLICY` entry specifies the name of the policy to use for user archives.

Table 2-212 BPARCHIVE_POLICY information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPARCHIVE_POLICY = policy_name</pre> By default, BPARCHIVE_POLICY is not present in the configuration file. By default, NetBackup uses the first policy that it finds that has the client and a user archive schedule. If it is used, this option should appear only once in the configuration file. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence, if it exists. See "Type of backup (schedule attribute)" on page 732.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

BPARCHIVE_SCHED option for NetBackup clients

This entry specifies the name of the schedule for user archives.

Table 2-213 BPARCHIVE_SCHED information

Usage	Description
Where to use	On NetBackup clients.

Table 2-213 BPARCHIVE_SCHED information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPARCHIVE_SCHED = schedule_name</pre> By default, BPARCHIVE_SCHED is not present in the configuration file. By default, NetBackup uses the first archive schedule in the first policy that it finds that contains this client. If it is used, this option should appear only once in the configuration file. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. See "Considerations for user schedules" on page 740.
Equivalent NetBackup web UI property	No equivalent exists.

BPBACKUP_POLICY option for NetBackup clients

This entry specifies the name of the policy to use for user backups.

Table 2-214 BPBACKUP_POLICY information

Usage	Description
Where to use	On NetBackup clients.

Table 2-214 BPBACKUP_POLICY information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPBACKUP_POLICY = policy_name</pre> By default, BPBACKUP_POLICY is not present in the configuration file. By default, NetBackup uses the first policy it finds that has both the client and a user backup schedule. If present, this option should appear only once in the configuration file. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. See "Type of backup (schedule attribute)" on page 732.
Equivalent NetBackup web UI property	No equivalent exists.

BPBACKUP_SCHED option for NetBackup clients

This entry specifies the name of the schedule to use for user backups.

Table 2-215 BPBACKUP_SCHED information

Usage	Description
Where to use	On NetBackup clients.

Table 2-215 BPBACKUP_SCHEDULE information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BPBACKUP_SCHEDULE = schedule_name</pre> By default, BPBACKUP_SCHEDULE is not present in the configuration file. By default, NetBackup uses the first policy it finds that contains both the client and a user backup schedule. If present, this option should appear only once in the configuration file. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. See "Considerations for user schedules" on page 740.
Equivalent NetBackup web UI property	No equivalent exists.

BUSY_FILE_ACTION option for NetBackup clients

The `BUSY_FILE_ACTION` entry directs the action that NetBackup performs on busy files when busy-file processing is enabled.

Table 2-216 BUSY_FILE_ACTION information

Usage	Description
Where to use	On NetBackup clients.

Table 2-216 BUSY_FILE_ACTION information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. <code>BUSY_FILE_ACTION = filename_template action_template</code> ■ filename_template The absolute pathname and file name of the busy file. The shell language metacharacters <code>*</code>, <code>?</code>, <code>[]</code>, <code>[-]</code> can be used for matching patterns of file names or parts of file names. ■ action_template Use one of the following parameters: ■ <code>MAIL mail</code> Directs NetBackup to email a busy file notification message to the user that the <code>BUSY_FILE_NOTIFY_USER</code> option specifies. ■ <code>REPEAT repeat [repeat_count]</code> Directs NetBackup to retry the backup on the specified busy file. A repeat count can be specified to control the number of backup attempts. The default repeat count is 1. ■ <code>IGNORE ignore</code> Directs NetBackup to exclude the busy file from processing. Multiple <code>BUSY_FILE_ACTION</code> entries are allowed. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Busy file settings. See “Busy file settings properties” on page 59.

BUSY_FILE_DIRECTORY option for NetBackup clients

The `BUSY_FILE_DIRECTORY` entry specifies the path to the busy-files working directory when busy-file processing is enabled.

Table 2-217 BUSY_FILE_DIRECTORY information

Usage	Description
Where to use	On NetBackup clients.

Table 2-217 BUSY_FILE_DIRECTORY information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BUSY_FILE_DIRECTORY = pathname</pre> By default, <code>BUSY_FILE_DIRECTORY</code> is not present in the configuration file. By default, NetBackup creates the <code>busy_files</code> directory in <code>/usr/opensv/netbackup</code>. If present, this option should appear only once in the configuration file. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence, if it exists.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Busy file settings. See "Busy file settings properties" on page 59.

BUSY_FILE_NOTIFY_USER option for NetBackup clients

The `BUSY_FILE_NOTIFY_USER` entry specifies who receives a notification when the `BUSY_FILE_ACTION` entry is set to `MAIL`.

Table 2-218 BUSY_FILE_NOTIFY_USER information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BUSY_FILE_NOTIFY_USER = email</pre> By default, <code>BUSY_FILE_NOTIFY_USER</code> is not present in the configuration file. By default, the email recipient is <code>root</code>. If present, this option should appear only once in the configuration file. The value in the user's <code>\$HOME/bp.conf</code> file takes precedence, if it exists.

Table 2-218 BUSY_FILE_NOTIFY_USER information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Busy file settings. See “Busy file settings properties” on page 59.

BUSY_FILE_PROCESSING option for NetBackup clients

The `BUSY_FILE_PROCESSING` entry lets the administrator control what NetBackup does when a file changes while it is in the process of being backed up.

Table 2-219 BUSY_FILE_PROCESSING information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>BUSY_FILE_PROCESSING = yes</pre> By default, <code>BUSY_FILE_PROCESSING</code> is not present in the configuration file and busy-file processing does not occur. If present, this option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Busy file settings. See “Busy file settings properties” on page 59.

CLIENT_NAME option for NetBackup clients

The `CLIENT_NAME` entry specifies the name of the client as it's known to NetBackup.

Table 2-220 CLIENT_NAME information

Usage	Description
Where to use	On NetBackup clients.

Table 2-220 CLIENT_NAME information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CLIENT_NAME = clientname</pre> This option should appear only once in the configuration file. If more than one entry appears, NetBackup observes only the last <code>CLIENT_NAME</code> entry that is listed for the option. The client name in a policy that backs up the client should match the client name that is specified with <code>CLIENT_NAME</code>. Note: Do not use an IP address as a client name in a policy or the backup may fail. Specify a host name instead. The <code>bp.conf</code> of the primary server does not require the addition of other clients, other than the primary server as <code>CLIENT_NAME = primary server name</code>. The name is added by default. During a restore, the default is to restore to the client that is named in the policy that was used for the backup. To redirect a restore to a different client, indicate the name of the target client in the Destination client field. To use the <code>bprestore</code> command, a parameter can be given to <code>bprestore</code> to indicate the destination client. The client installation procedures automatically set <code>CLIENT_NAME</code> to the value that is specified in <code>ftp_to_client</code> command or <code>install_client</code> command in the installation scripts. If the value is not in any <code>bp.conf</code> file, NetBackup uses the value that the <code>gethostname()</code> library function returns. See “Client name properties” on page 63.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Client name.

COMPRESS_SUFFIX option for NetBackup clients

The `COMPRESS_SUFFIX` entry specifies a list of file extensions. During a backup, NetBackup does not compress files with these extensions because the file may already be in a compressed format.

Table 2-221 COMPRESS_SUFFIX information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>COMPRESS_SUFFIX = .suffix</pre> By default, <code>COMPRESS_SUFFIX</code> is not present in the <code>bp.conf</code> file. This option has a reasonable default. Change only if problems result. Multiple <code>COMPRESS_SUFFIX</code> entries are allowed. Do not use wildcards to specify these extensions. Do not specify <code>.A*</code> or <code>.A [1-9]</code> (For example, specify <code>.A1</code>)
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > UNIX client > Client settings > Do not compress files with these file extensions. See “Client settings properties for UNIX clients” on page 73.

CRYPT_CIPHER option for NetBackup clients

The `CRYPT_CIPHER` entry applies to clients with the NetBackup Encryption option installed.

Table 2-222 CRYPT_CIPHER information

Usage	Description
Where to use	On NetBackup clients.

Table 2-222 CRYPT_CIPHER information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CRYPT_CIPHER = value</pre> Where <i>value</i> is one of the following: ■ AES-128-CFB (used when no method is specified; default) ■ AES-256-CFB ■ BF-CFB ■ DES-EDE-CFB This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Encryption. See “Encryption properties” on page 88. For information about NetBackup encryption, see the NetBackup Security and Encryption Guide.

CRYPT_KIND option for NetBackup clients

The `CRYPT_KIND` entry on the client determines whether the standard encryption or legacy encryption is used in the backup. Normally, `CRYPT_KIND` is set automatically.

The `CRYPT_KIND` entry applies to clients with the NetBackup Encryption option installed.

See the [NetBackup Security and Encryption Guide](#).

Table 2-223 CRYPT_KIND information

Usage	Description
Where to use	On NetBackup clients.

Table 2-223 CRYPT_KIND information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CRYPT_KIND = type</pre> The following values can be entered: ■ NONE No encryption is used on the client (default) ■ LEGACY Legacy pertains to 40-bit and 56-bit data encryption standard (DES). Legacy encryption is not recommended. ■ STANDARD Standard pertains to 128-bit and 256-bit encryption (AES, 3DES, Blowfish cipher). Standard encryption is recommended. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Encryption. See “Encryption properties” on page 88.

CRYPT_OPTION option for NetBackup clients

The `CRYPT_OPTION` entry specifies the encryption options on NetBackup clients. NetBackup creates this entry automatically in the `/usr/openv/netbackup/bp.conf` file on a UNIX client when the `bpinst_crypt` command is run on the NetBackup primary server.

The `CRYPT_OPTION` entry applies to clients with the NetBackup Encryption option installed.

For information about these commands, see the [NetBackup Commands Reference Guide](#).

Do not alter the entry or create this file manually unless it was accidentally deleted.

Table 2-224 CRYPT_OPTION information

Usage	Description
Where to use	On NetBackup clients.

Table 2-224 CRYPT_OPTION information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CRYPT_OPTION = type</pre> The following values can be entered: ■ <code>DENIED denied</code> Specifies that the client does not permit encrypted backups. If the server requests an encrypted backup, it is considered an error. This option is the default for a client that has not been configured for encryption. ■ <code>ALLOWED allowed</code> Specifies that the client allows either encrypted or unencrypted backups. ■ <code>REQUIRED required</code> Specifies that the client requires encrypted backups. If this value is specified and the server requests an unencrypted backup, it is considered an error. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Encryption. See "Encryption properties" on page 88.

CRYPT_STRENGTH option for NetBackup clients

The `CRYPT_STRENGTH` entry specifies the encryption strength on NetBackup clients. It applies to clients with the NetBackup Encryption option installed.

See the [NetBackup Security and Encryption Guide](#).

Table 2-225 CRYPT_STRENGTH information

Usage	Description
Where to use	On NetBackup clients.

Table 2-225 CRYPT_STRENGTH information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CRYPT_STRENGTH = value</pre> The allowable values are as follows: ■ <code>DES_40 des_40</code> Specifies 40-bit DES encryption. 40-bit is the default value for a client that has not been configured for encryption. ■ <code>DES_56 des_56</code> Specifies 56-bit DES encryption. NetBackup creates this entry automatically on a UNIX client when the <code>bpinst_crypt</code> command is run on the NetBackup primary server. This entry should appear only once in the configuration file. Do not alter the entry or create it manually unless it was accidentally deleted.
Equivalent NetBackup web UI property	No equivalent exists. For standard encryption options, see the following topic. See “Encryption properties” on page 88.

CRYPT_LIBPATH option for NetBackup clients

The `CRYPT_LIBPATH` entry specifies the directory that contains the encryption libraries for NetBackup clients. NetBackup creates this entry automatically in the `/usr/opensv/netbackup/bp.conf` file on a UNIX client when the `bpinst_crypt` command is run on the NetBackup primary server.

The `CRYPT_LIBPATH` entry applies to clients with the NetBackup Encryption option installed. Do not alter the entry or create it manually unless it was accidentally deleted. This entry should appear only once in the configuration file.

See the [NetBackup Security and Encryption Guide](#).

Table 2-226 CRYPT_LIBPATH information

Usage	Description
Where to use	On NetBackup clients.

Table 2-226 CRYPT_LIBPATH information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CRYPT_LIBPATH = directory</pre> If necessary, create the entry in the following locations: ■ The default value on Windows systems is <code>install_path\NetBackup\bin\</code>. Where <code>install_path</code> is the directory where NetBackup is installed and by default is <code>C:\Program Files\Cohesity NetBackup</code>. ■ The default value on UNIX systems is <code>/usr/opensv/lib/</code>
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client properties > Encryption. See “Encryption properties” on page 88.

CRYPT_KEYFILE option for NetBackup clients

The `CRYPT_KEYFILE` entry specifies the file that contains the encryption keys on NetBackup clients. NetBackup creates this entry automatically in the `/usr/opensv/netbackup/bp.conf` file on a UNIX client when the `bpinst_crypt` command is run on the NetBackup primary server.

See the [NetBackup Security and Encryption Guide](#).

Table 2-227 CRYPT_KEYFILE information

Usage	Description
Where to use	On NetBackup clients.

Table 2-227 CRYPT_KEYFILE information (*continued*)

Usage	Description
How to use	The <code>CRYPT_KEYFILE</code> entry applies to clients with the NetBackup Encryption option installed. Do not alter the entry or create it manually unless it was accidentally deleted. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>CRYPT_KEYFILE = directory</pre> The default values follow: - On Windows systems <code>install_path\bin\keyfile.dat</code> Where <code>install_path</code> is the directory where NetBackup is installed and by default is <code>C:\Program Files\Cohesity NetBackup</code>. - On UNIX systems <code>/usr/opensv/netbackup/keyfile</code> This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists. For standard encryption options, see the following topic. See “Encryption properties” on page 88.

DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients

The `DO_NOT_RESET_FILE_ACCESS_TIME` entry specifies that if a file is backed up, its access time (`atime`) displays the time of the backup. The default is that NetBackup preserves the access time by resetting it to the value it had before the backup.

Table 2-228 DO_NOT_RESET_FILE_ACCESS_TIME information

Usage	Description
Where to use	On NetBackup clients.

Table 2-228 DO_NOT_RESET_FILE_ACCESS_TIME information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DO_NOT_RESET_FILE_ACCESS_TIME</pre> This option should appear only once in the configuration file. Note: <code>DO_NOT_RESET_FILE_ACCESS_TIME</code> must be enabled if <code>USE_CTIME_FOR_INCREMENTALS</code> is enabled. Setting these options causes the file <code>atime</code> to be updated every time they are backed up. When the <code>atime</code> updates, it appears as if the files have been recently used.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client UNIX client > Client settings > Reset file access time to the value before backup. Note: The property label/description is the opposite of the configuration setting. Thus when the property is enabled, <code>DO_NOT_RESET_FILE_ACCESS_TIME</code> is disabled. Similarly when the property is disabled, <code>DO_NOT_RESET_FILE_ACCESS_TIME</code> is enabled. This property must be left disabled when <code>USE_CTIME_FOR_INCREMENTALS</code> is enabled. See “Client settings properties for UNIX clients” on page 73.

DTE_CLIENT_MODE for clients

The `DTE_CLIENT_MODE` option specifies the data-in-transit encryption (DTE) mode that is set on the NetBackup client.

Table 2-229 DTE_CLIENT_MODE information

Usage	Description
Where to use	On NetBackup clients.

Table 2-229 DTE_CLIENT_MODE information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>DTE_CLIENT_MODE = AUTOMATIC ON OFF</pre> By default, the DTE mode for 9.1 clients is set to <code>OFF</code> and for 10.0 and later clients, it is set to <code>AUTOMATIC</code>. ■ If the <code>DTE_CLIENT_MODE</code> option is set to <code>AUTOMATIC</code>, the client follows the DTE mode that is set at the global level: <code>Enforced</code>, <code>Preferred On</code>, or <code>Preferred Off</code>. ■ If the option is set to <code>ON</code>, data-in-transit encryption is enabled for the client. ■ If the option is set to <code>OFF</code>, data-in-transit encryption is disabled for the client. This setting can be used to exclude a client for encryption if the global DTE mode is set to <code>Preferred On</code>. Note: If the global DTE mode is set to <code>Enforced</code>, jobs fail for the NetBackup clients that have the <code>DTE_CLIENT_MODE</code> option set to 'OFF' and also for the hosts earlier than 9.1.
Equivalent NetBackup web UI property	No equivalent exists. Global settings are configured in Settings > Global security > Secure communication > Data-in-transit encryption.

ENABLE_DATA_CHANNEL_ENCRYPTION for clients

The `ENABLE_DATA_CHANNEL_ENCRYPTION` option specifies if the data channel encryption is enabled for communication with NAT clients and NAT servers (or NAT hosts).

If a NAT host is configured in your NetBackup domain, data channel encryption is enabled by default.

Table 2-230 ENABLE_DATA_CHANNEL_ENCRYPTION information

Usage	Description
Where to use	On NetBackup clients.

Table 2-230 ENABLE_DATA_CHANNEL_ENCRYPTION information
(continued)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To disable data channel encryption, use the following format: <pre>ENABLE_DATA_CHANNEL_ENCRYPTION = FALSE</pre>
Equivalent NetBackup web UI property	No equivalent exists in host properties.

IGNORE_XATTR option for NetBackup clients

By default, extended attribute files (Solaris 9 or later) and named data streams are backed up. Use `IGNORE_XATTR` to exclude extended attributes and named data streams from backups. (`IGNORE_XATTR` was formerly `IGNORE_XATTR_SOLARIS`.)

Table 2-231 IGNORE_XATTR information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>IGNORE_XATTR</pre> NetBackup does not check for extended attributes or named data streams if the <code>IGNORE_XATTR</code> entry is present in the <code>bp.conf</code> file. This option should appear only once in the configuration file. See “About backing up and restoring extended attribute files and named data streams” on page 807.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

Backing up and restoring the ACLs, extended attributes, and metadata attributes for GPFS volumes

NetBackup recognizes the access control lists, extended attributes, and metadata attributes of General Parallel File System files and folders. By default, NetBackup backs up and restores all of these for GPFS volumes. No additional configuration is necessary.

Restoring files with GPFS Extended Attributes to operating systems other than AIX or RHEL or file systems other than GPFS may generate errors or failed jobs. The errors occur because the restore target does not recognize the metadata.

Preventing the backup or restore of the ACL and extended attributes for a GPFS volume

- To prevent backups of GPFS extended attributes:

Add the `IGNORE_XATTR` entry to the `bp.conf` file on the client:

```
/usr/opensv/netbackup/bp.conf
```

The entry does not require a value setting; the entry is only `IGNORE_XATTR`.

- To prevent backups of the GPFS ACL:

Add the `IGNORE_ACL` touch file to the client:

```
/usr/opensv/netbackup/IGNORE_ACL
```

Note: If extended attributes are allowed to be backed up, the ACL is also backed up, regardless of whether or not the `IGNORE_ACL` touch file is present.

- To prevent restores of GPFS extended attributes (provided that they were backed up in the first place):

Add the `IGNORE_XATTR` touch file to the client:

```
/usr/opensv/netbackup/IGNORE_XATTR
```

- To prevent restores of the GPFS ACL (provided that it was backed up in the first place):

Add the `IGNORE_ACL` touch file to the client:

```
/usr/opensv/netbackup/IGNORE_ACL
```

Note: If extended attributes are allowed for restore, the ACL is also restored, regardless of whether or not the `IGNORE_ACL` touch file is present.

Considerations when backing up and restoring GPFS extended attributes

- Upon restore of a file or folder that has an ACL or extended attributes, the ACL and extended attributes are also restored if they were originally backed up. However, if the restore job was interrupted, the last file that was in-process may be skipped when the restore job resumes. In that case, the ACL and the extended attributes of that file and all subsequent files will not be restored.
- The `IGNORE_XATTR` option instructs NetBackup to ignore extended attributes when backing up files and folders that have extended attributes. In case of a NetBackup Accelerator initial backup, if `IGNORE_XATTR` is added to the `bp.conf` file or the registry and then is later removed for subsequent backups, perform a backup with the **Accelerator forced rescan** option enabled.
See [“Accelerator forced rescan option \(schedule attribute\)”](#) on page 742.
- If the ACL interface is used to change ACL permissions after a backup is performed, the restore may not preserve the ACL upon restore.

Support for GPFS metadata attributes: storage pools, metadata replication, and data replication

NetBackup has added support for the backup and restore of the following GPFS metadata attributes: storage pools, metadata replication, and data replication. By default, NetBackup backs up and restores the ACLs, the extended attributes, and these additional metadata attributes for GPFS volumes. No additional configuration is necessary.

The following information is pertinent if, after restoring a file, the `illplaced` or the `illreplicated` flags display in the status of the file. (View the status of the file by using the GPFS `mmfsattr` command.)

The `illplaced` flag displays if the file was restored to a different GPFS storage pool than where it was when it was backed up. The change could be due to a change in creation rules or the result of migration rules in the GPFS file placement policy. Or the GPFS administrator may have used the `mmchattr` command to manually move the file to a different storage pool.

The `illreplicated` flag displays when the storage pool containing the restored file has fewer failure groups than required by the data replication setting or the system pool has fewer failure groups than the metadata replication setting.

The GPFS administrator can rebalance the replication factor of the file and resolve the `illplaced` or `illreplicated` flags by running one of two GPFS commands, depending on the number of the files that are involved and the network traffic:

- For a single file: `mmrestripefile`

- For the complete file system: `mmrestripefs`

Note: Restoring files with GPFS attributes and ACLs to an alternate platform does not restore the metadata attributes or the ACLs. The restore may generate an error such as “Invalid system call.” (Extended attributes can be restored to the alternate platform, however.)

INFORMIX_HOME option for NetBackup clients

The `INFORMIX_HOME` entry specifies the path to the Informix home directory and is required when the client uses NetBackup for Informix.

Table 2-232 `INFORMIX_HOME` information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>INFORMIX_HOME</pre> This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

KEEP_DATABASE_COMM_FILE option for NetBackup clients

The `KEEP_DATABASE_COMM_FILE` entry causes NetBackup to keep database agent logs for seven days. The default is that NetBackup keeps database agent logs for only one day.

Table 2-233 `KEEP_DATABASE_COMM_FILE` information

Usage	Description
Where to use	On NetBackup clients.

Table 2-233 KEEP_DATABASE_COMM_FILE information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>KEEP_DATABASE_COMM_FILE</pre> For example, add it to a client that runs NetBackup for Informix. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

KEEP_LOGS_DAYS option for NetBackup clients

The `KEEP_LOGS_DAYS` entry specifies how long to keep job and progress logs for **Backup, Archive, and Restore**. The default is 3 days.

NetBackup writes these files in the following directories:

- `/usr/opensv/netbackup/logs/user_ops/username/jobs`
- `/usr/opensv/netbackup/logs/user_ops/username/logs`

A directory exists for each user that uses the **Backup, Archive, and Restore** console.

For NetBackup Java capable platforms, the NetBackup Java interface log files are written in a subdirectory that is called `nbjlogs`. The `KEEP_LOGS_DAYS` option does not remove log files from this directory. The administrator can clean up these logs in this directory according to the requirements of their organization.

Table 2-234 KEEP_LOGS_DAYS information

Usage	Description
Where to use	On NetBackup clients.

Table 2-234 KEEP_LOGS_DAYS information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option file. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>KEEP_LOGS_DAYS</pre> This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Windows Client > Client settings > Keep status of user-directed backups, archives and restores Hosts > Host properties > Select the client > UNIX Client > Client settings > Keep status of user-directed backups, archives and restores

LIST_FILES_TIMEOUT option for NetBackup clients

The `LIST_FILES_TIMEOUT` entry specifies how long to wait for a response from the server when it lists files by using the client-user interface or `bplist`. If this time is exceeded, the user receives a `socket read failed` error even if the server continues to process the user's request. The default is that `LIST_FILES_TIMEOUT` is not in any `bp.conf` file and NetBackup uses a value of 30 minutes.

Table 2-235 LIST_FILES_TIMEOUT information

Usage	Description
Where to use	On NetBackup clients.
How to use	For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>LIST_FILES_TIMEOUT</pre> This option should appear only once in the configuration file. The value in the <code>\$HOME/bp.conf</code> file takes precedence if it exists.
Equivalent web UI property	Hosts > Host properties > Select the client > Edit client > Timeouts. See "Timeouts properties" on page 175.

LOCKED_FILE_ACTION option for NetBackup clients

The `LOCKED_FILE_ACTION` entry specifies the behavior of NetBackup when it backs up a file that has mandatory file locking enabled in its file mode. (See `chmod(1)`). If this entry is set to `SKIP`, NetBackup skips the files that currently have mandatory locking set by another process. NetBackup logs a message to this effect.

Table 2-236 LOCKED_FILE_ACTION information

Usage	Description
Where to use	On NetBackup UNIX/Linux clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>LOCKED_FILE_ACTION = SKIP</pre> This option should appear only once in the configuration file. Note: <code>SKIP</code> is the only legal value for this entry. The default is that NetBackup waits for files to become unlocked.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > UNIX client > Client settings. See “Client settings properties for UNIX clients” on page 73.

MEDIA_SERVER option for NetBackup clients

The `MEDIA_SERVER` entry specifies that the listed computer is a media server only. Computers that are listed as media servers can back up and restore clients, but have limited administrative privileges.

Table 2-237 MEDIA_SERVER information

Usage	Description
Where to use	On NetBackup clients.

Table 2-237 MEDIA_SERVER information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MEDIA_SERVER = media_server_name</pre> This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Servers. See “SERVER option for NetBackup servers” on page 307.

MEGABYTES_OF_MEMORY option for NetBackup clients

The `MEGABYTES_OF_MEMORY` entry specifies how much memory is available on the client to use to compress files during backup. If compression is selected, the client software uses this value to determine how much space to request for the compression tables. The more memory that is available to the compress code, the greater the compression. The percentage of computer resources that are used is also greater. If other processes also need memory, use a maximum value of one half the actual physical memory on a computer to avoid excessive swapping.

Table 2-238 MEGABYTES_OF_MEMORY information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>MEGABYTES_OF_MEMORY = memory_value</pre> The default is that NetBackup assumes a value of zero megabytes. This option should appear only once in the configuration file. Note: The <code>MEGABYTES_OF_MEMORY</code> entry has a reasonable default. Change it only if problems are encountered.

Table 2-238 MEGABYTES_OF_MEMORY information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > UNIX client settings. See “Client settings properties for UNIX clients” on page 73.

MSSQL_CONFIG_LIST for SQL Server clients

The `MSSQL_CONFIG_LIST` option provides support for the non-readable secondary instance that is hidden. No configuration on the secondary that provides the port number of the primary to NetBackup, so the NetBackup user must provide it.

Table 2-239 MSSQL_CONFIG_LIST information

Usage	Description
Where to use	On NetBackup SQL Server clients.
How to use	Use the <code>bpgetconfig</code> and the <code>bpsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To add the port number of the primary, use the following format: <pre>MSSQL_CONFIG_LIST = "hostname\instance,port"</pre> For example: <pre>MSSQL_CONFIG_LIST = "myhost\SQL2K22,1633".</pre>
Equivalent NetBackup web UI property	No equivalent exists in host properties. A user with the necessary RBAC permissions can also use the <code>hostProperties</code> API endpoint to configure this setting. See the NetBackup for Microsoft SQL Server Administrator's Guide for details.

To configure the port settings on multiple nodes

- 1 Set the `host\instance,port` on the first node.

```
C:\Program Files\Cohesity NetBackup\NetBackup\bin\admincmd>bpsetconfig -h
bpsetconfig> MSSQL_CONFIG_LIST=host16vm5\SQL2K22,1633
bpsetconfig> MSSQL_CONFIG_LIST=host16vm6\SQL2K22,1634
bpsetconfig> ^Z
```

- 2 Use `bpgetconfig` to copy the settings to a file.

```
F:\Program Files\Cohesity NetBackup\NetBackup\bin\admincmd>bpgetconfig -M
findstr MSSQL_CONFIG_LIST > settings.out
```

- 3 Use the file to set the configuration on all the other nodes.

```
C:\Program Files\Cohesity NetBackup\NetBackup\bin\admincmd>bpsetconfig -h
```

- 4 Use `bpgetconfig` to validate the setting.

```
F:\Program Files\Cohesity NetBackup\NetBackup\bin\admincmd>bpgetconfig -M
findstr MSSQL_CONFIG_LIST
MSSQL_CONFIG_LIST = host16vm5\SQL2K22,1633
MSSQL_CONFIG_LIST = host16vm6\SQL2K22,1634
```

MSSQL_ODBC_ENCRYPT_CONNECTION for SQL Server clients

The `MSSQL_ODBC_ENCRYPT_CONNECTION` option determines whether to encrypt the connection to a target SQL Server client using TLS. For NetBackup SQL Server clients that are updated to 10.4 or later, the SQL Server ODBC connections from the client to a target SQL Server instance are encrypted by default.

Table 2-240 MSSQL_ODBC_ENCRYPT_CONNECTION information

Usage	Description
Where to use	On NetBackup SQL Server clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To disable data channel encryption, use the following format: <pre>MSSQL_ODBC_ENCRYPT_CONNECTION = NO</pre>

Table 2-240 MSSQL_ODBC_ENCRYPT_CONNECTION information
(continued)

Usage	Description
Equivalent NetBackup web UI property	No equivalent exists in host properties. A user with the necessary RBAC permissions can also use the <code>hostProperties</code> API endpoint to configure this setting. See the NetBackup for Microsoft SQL Server Administrator's Guide for details.

More information

See “[MSSQL_ODBC_TRUST_SERVER_CERTIFICATE for SQL Server clients](#)” on page 361.

See “[MSSQL_ODBC_PREFERRED_DRIVER for SQL Server clients](#)” on page 360.

MSSQL_ODBC_PREFERRED_DRIVER for SQL Server clients

The name of the supported SQL Server ODBC driver to use during the connection.

Table 2-241 MSSQL_ODBC_PREFERRED_DRIVER information

Usage	Description
Where to use	On NetBackup SQL Server clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. The default value is <code>OLDEST</code>. Customers that have strict security policies and concerns should use the <code>NEWEST</code> driver or whatever version their company has certified. To automatically select the oldest or newest driver, use the following format: <pre>MSSQL_ODBC_PREFERRED_DRIVER = OLDEST NEWEST</pre> For specific driver names, create a separate entry for each driver name. <pre>MSSQL_ODBC_PREFERRED_DRIVER = ODBC Driver 18 for SQL Server MSSQL_ODBC_PREFERRED_DRIVER = ODBC Driver 17 for SQL Server</pre>

Table 2-241 MSSQL_ODBC_PREFERRED_DRIVER information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	No equivalent exists in host properties. A user with the necessary RBAC permissions can also use the <code>hostProperties</code> API endpoint to configure this setting. See the NetBackup for Microsoft SQL Server Administrator's Guide for details.

More information

See “[MSSQL_ODBC_ENCRYPT_CONNECTION for SQL Server clients](#)” on page 359.

See “[MSSQL_ODBC_TRUST_SERVER_CERTIFICATE for SQL Server clients](#)” on page 361.

MSSQL_ODBC_TRUST_SERVER_CERTIFICATE for SQL Server clients

The `MSSQL_ODBC_TRUST_SERVER_CERTIFICATE` option determines whether to trust the target SQL Server instance's certificate.

For 10.3.1 and later clients, `MSSQL_ODBC_TRUST_SERVER_CERTIFICATE` is set to `NO` by default to prevent unexpected connection failures on upgrade.

Table 2-242 MSSQL_ODBC_TRUST_SERVER_CERTIFICATE information

Usage	Description
Where to use	On NetBackup SQL Server clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. To disable trust of the target instance's certificate, use the following format: <pre>MSSQL_ODBC_TRUST_SERVER_CERTIFICATE = NO</pre>
Equivalent NetBackup web UI property	No equivalent exists in host properties. A user with the necessary RBAC permissions can also use the <code>hostProperties</code> API endpoint to configure this setting. See the NetBackup for Microsoft SQL Server Administrator's Guide for details.

More information

See “[MSSQL_ODBC_ENCRYPT_CONNECTION for SQL Server clients](#)” on page 359.

See “[MSSQL_ODBC_TRUST_SERVER_CERTIFICATE for SQL Server clients](#)” on page 361.

NFS_ACCESS_TIMEOUT option for NetBackup clients

The `NFS_ACCESS_TIMEOUT` entry specifies the number of seconds that the backup process waits to process an NFS mount table. After the time is exceeded, the process considers an NFS file system to be unavailable.

Table 2-243 NFS_ACCESS_TIMEOUT information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>NFS_ACCESS_TIMEOUT = <i>seconds</i></pre> By default, <code>NFS_ACCESS_TIMEOUT</code> is set to five seconds. This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists.

OLD_VNETD_CALLBACK option for NetBackup clients

The `OLD_VNETD_CALLBACK` entry specifies that the client should use the client-direct restore path rather than use the `bptm` process on a NetBackup media server. The client-direct restore path bypasses the NetBackup media server for the restore jobs. Configure this entry on the NetBackup clients that you want to receive data directly from a **Media Server Deduplication Pool** storage server.

Before you decommission a media server that hosts a **Media Server Deduplication Pool**, deactivate MSDP on that media server. See the [NetBackup Deduplication Guide](#) for your release.

Table 2-244 OLD_VNETD_CALLBACK information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>OLD_VNETD_CALLBACK = YES</pre> This entry should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

REPORT_CLIENT_DISCOVERIES option for NetBackup clients

By default, the NetBackup Discovery Service (`nbdisco`) runs on all clients in the NetBackup environment. The service helps you build Intelligent Policies by reporting to the primary server when it finds instances of applications (such as Oracle).

Table 2-245 REPORT_CLIENT_DISCOVERIES information

Usage	Description
Where to use	On NetBackup clients.

Table 2-245 REPORT_CLIENT_DISCOVERIES information (*continued*)

Usage	Description
How to use	By default, <code>REPORT_CLIENT_DISCOVERIES</code> is not present in the configuration file. When <code>REPORT_CLIENT_DISCOVERIES</code> is not present, the NetBackup Discovery Service is enabled. Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>REPORT_CLIENT_DISCOVERIES = TRUE FALSE</pre> This option should appear only once in the configuration file. ■ If the entry is set to <code>FALSE</code>, the <code>nbdisco</code> process on that client stops reporting discoveries to the primary server. The service shuts down within 10 minutes after being set to <code>FALSE</code> and remains down. ■ To turn on the Discovery Service again, either change the entry on that client to <code>REPORT_CLIENT_DISCOVERIES = TRUE</code> or remove the entire option. Then, run <code>bp.start_all</code> on the client to restart the service. ■ To set this value on a client remotely, run the following command from the primary server: <pre>nbsetconfig -h clientname</pre>
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

RESTORE_RETRIES option for NetBackup clients

The `RESTORE_RETRIES` entry specifies the number of times to retry a restore after a failure.

Table 2-246 RESTORE_RETRIES information

Usage	Description
Where to use	On NetBackup clients.

Table 2-246 RESTORE_RETRIES information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RESTORE_RETRIES = number_of_retries</pre> The default is 0 (no retries). This option should appear only once in the configuration file. Note: The <code>RESTORE_RETRIES</code> entry has a reasonable default. Change it only if problems are encountered.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Universal settings > Restore retries. See “Universal settings properties” on page 178.

RMAN_OUTPUT_DIR for NetBackup clients

The `RMAN_OUTPUT_DIR` specifies which directory to place the RMAN input and output locally on the client for Oracle Intelligent Policy backups. NetBackup does not clean up the log files so the Oracle user has to clean up the log files manually. The log is only created when a backup is run using an Oracle Intelligent Policy. Only one `RMAN_OUTPUT_DIR` entry per client is allowed in a Windows environment. In a UNIX environment, each user can place the output in a different location by adding the `RMAN_OUTPUT_DIR` entry to `$HOME/bp.conf` file. The value in the `$HOME/bp.conf` file takes precedence if it exists.

The following are examples of `RMAN_OUTPUT_DIR` entries:

Windows: `install_path\oracle\oracle_logs\RMAN`

UNIX: `/oracle/oracle_logs/rman`

For information about `RMAN_OUTPUT_DIR`, see the [NetBackup for Oracle Administrator's Guide](#).

Table 2-247 RMAN_OUTPUT_DIR information

Usage	Description
Where to use	On NetBackup clients.

Table 2-247 RMAN_OUTPUT_DIR information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>RMAN_OUTPUT_DIR = directory_name</pre> The <i>directory_name</i> is a directory to which the Oracle user has permission to write.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

SERVER option for NetBackup clients

The `SERVER` entry defines the list of NetBackup primary servers and media servers that can access the NetBackup client. During client installation, `SERVER` is automatically set to the name of the primary primary server for this client.

The `SERVER` entries must be added for other primary servers and for media servers for this client. The client needs to have certificates from all the primary servers to communicate with the server. To get the certificate, the client should have entries of all the connected primary servers in the client `bp.conf` file.

Table 2-248 SERVER information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Note: Every <code>SERVER</code> entry in a client <code>bp.conf</code> file must be a NetBackup primary or media server. That is, each system that is listed as a <code>SERVER</code> must have either NetBackup primary or media server software installed. The client service on some clients cannot be started if the client name is incorrectly listed as a server. If you configure media servers, each media server must have a <code>SERVER</code> or a <code>MEDIA_SERVER</code> entry in the <code>bp.conf</code> file of the client.

Table 2-248 SERVER information (*continued*)

Usage	Description
Example	The following is an example <code>bp.conf</code> file on a client: <pre>SERVER = Primary_server (default primary server) SERVER = NB_server (other primary server) SERVER = Media_server_#1 MEDIA_SERVER = Media_server_#2 . . .</pre> The first <code>SERVER</code> entry denotes the primary server to which the client connects to by default for any requests. (For example, to back up, to list, or to restore files). The <code>SERVER</code> entry must be present in the <code>/usr/opensv/netbackup/bp.conf</code> file on all UNIX clients. The <code>SERVER</code> entry is the only required entry in the <code>bp.conf</code> file for clients. The <code>SERVER</code> entry is not used in a <code>\$HOME/bp.conf</code> file. On NetBackup UNIX servers, the <code>SERVER</code> entry applies to both client and the server. See “RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers” on page 304. See “MEDIA_SERVER option for NetBackup clients” on page 356.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Servers. See “SERVER option for NetBackup servers” on page 307.

SUBSCRIBER_HEARTBEAT_TIMEOUT for clients

The `SUBSCRIBER_HEARTBEAT_TIMEOUT` option enables the subscriber service on a NAT client or NAT server to send heartbeats (or signals) for communication with the primary server.

The option value is defined in seconds.

Table 2-249 SUBSCRIBER_HEARTBEAT_TIMEOUT information

Usage	Description
Where to use	On NetBackup clients.

Table 2-249 SUBSCRIBER_HEARTBEAT_TIMEOUT information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. If the option is set to a value other than zero, the subscriber service sends heartbeats. To disable the heartbeats, use the following format: <pre>SUBSCRIBER_HEARTBEAT_TIMEOUT = 0</pre>
Equivalent NetBackup web UI property	No equivalent exists.

SYBASE_HOME option for NetBackup clients

The `SYBASE_HOME` entry specifies the path to the Sybase home directory. The entry is required for NetBackup to use Sybase to back up Sybase databases.

Table 2-250 SYBASE_HOME information

Usage	Description
Where to use	On a NetBackup for Sybase client.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>SYBASE_HOME = path_to_Sybase_home_directory</pre> The default is that <code>SYBASE_HOME</code> is not in the configuration file. This option should appear only once in the configuration file. For information about these commands, see the NetBackup for Sybase Administrator's Guide.
Equivalent NetBackup web UI property	No equivalent exists in the host properties.

USE_CTIME_FOR_INCREMENTALS option for NetBackup clients

The `USE_CTIME_FOR_INCREMENTALS` entry changes how NetBackup determines whether or not a file has changed. This entry causes the client software to use both modification time and inode change time during incremental backups to determine if a file has changed. (`mtime` and `ctime`.)

Table 2-251 `USE_CTIME_FOR_INCREMENTALS` information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>USE_CTIME_FOR_INCREMENTALS</pre> This option should appear only once in the configuration file. By default, NetBackup uses only <code>mtime</code>. Note: If you specify <code>USE_CTIME_FOR_INCREMENTALS</code>, you must also specify <code>DO_NOT_RESET_FILE_ACCESS_TIME</code>. Setting these options causes the file <code>atime</code> to be updated every time they are backed up. When the <code>atime</code> updates, it appears as if the files have been recently used. See “DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients” on page 347.
Equivalent NetBackup web UI property	No equivalent exists.

USE_FILE_CHG_LOG option for NetBackup clients

The `USE_FILE_CHG_LOG` entry specifies whether NetBackup uses the file change log on VxFS clients. The default is off.

Table 2-252 `USE_FILE_CHG_LOG` information

Usage	Description
Where to use	On NetBackup clients.

Table 2-252 USE_FILE_CHG_LOG information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>USE_FILE_CHG_LOG</pre> This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client UNIX client > Client settings > Use VxFS File Change Log for incremental backups. See “VxFS file change log (FCL) for incremental backups property” on page 75.

USEMAIL option for NetBackup clients

The `USEMAIL` entry specifies the email address where NetBackup sends status on the outcome of operations for a UNIX client.

Table 2-253 USEMAIL information

Usage	Description
Where to use	On NetBackup UNIX clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>USEMAIL = name1@xxx.com,name2@xxx.com,name3@xxx.com</pre> Note: Separate multiple email addresses using a comma, with no spaces. Add <code>USEMAIL</code> as follows: ■ If the <code>USEMAIL</code> specifies an address, NetBackup sends automatic backup and manual backup status to that address. ■ If the <code>\$HOME/bp.conf</code> file specifies an address, NetBackup also sends status on the success or failure of user operations to that address. This option should appear only once in the configuration file.

Table 2-253 USEMAIL information (*continued*)

Usage	Description
Equivalent NetBackup web UI property	Hosts > Host properties > Select the client > Edit client > Universal settings. See " Universal settings properties " on page 178.

WEB_SERVER_TUNNEL option for NetBackup clients

Add the WEB_SERVER_TUNNEL option to use a specific media server that creates the HTTP tunnel for connecting to a specific primary server. This option overrides the media and primary server list that is automatically generated on the client for sending web service connection requests.

For more information, refer to the **About the communication between a NetBackup client located in a demilitarized zone and a primary server through an HTTP tunnel** section in the *NetBackup Security and Encryption Guide*.

Table 2-254 WEB_SERVER_TUNNEL information

Usage	Description
Where to use	On NetBackup clients.

Table 2-254 WEB_SERVER_TUNNEL information (*continued*)

Usage	Description
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>WEB_SERVER_TUNNEL = <primary> <media></pre> The option uses the following parameters: ■ <i>primary</i> is the hostname or IP address or FQDN of the primary server that should be the recipient of the web service connection requests. ■ <i>media</i> is the hostname or IP address or FQDN of the media server that sets up the connection via an HTTP tunnel. For a multi-domain setup, you can add multiple entries on separate lines. These entries can include a single primary server and different media servers or different identities of the media servers like IP addresses, host names, and Fully Qualified Domain Names (FQDN). For example, <pre>WEB_SERVER_TUNNEL=<primary> <media> WEB_SERVER_TUNNEL=<primary1> <media1> WEB_SERVER_TUNNEL=<primary1> <media1_IP address></pre> By default, the <code>WEB_SERVER_TUNNEL</code> option is not present in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists.

WEB_SERVER_TUNNEL_USE option for NetBackup clients

In a demilitarized zone (DMZ), the client uses a sequence of steps to set up communication with the primary server. By default, the client tries to communicate with the primary directly, and if the connection fails, tries to communicate using an HTTP tunnel on the media server. You can use the `WEB_SERVER_TUNNEL_USE` option to change the default behavior.

For more information, refer to the **About the communication between a NetBackup client located in a demilitarized zone and a primary server through an HTTP tunnel** section in the *NetBackup Security and Encryption Guide*.

Table 2-255 WEB_SERVER_TUNNEL_USE information

Usage	Description
Where to use	On NetBackup clients.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option. For information about these commands, see the NetBackup Commands Reference Guide. Use the following format: <pre>WEB_SERVER_TUNNEL_USE = AUTO ALWAYS NEVER</pre> The option uses the following parameters: ■ <code>AUTO</code> is the default value that uses an auto-routing algorithm. ■ <code>ALWAYS</code> defines that the connection should always use an HTTP tunnel. You can use this option for NetBackup clients that are in a DMZ. ■ <code>NEVER</code> defines that the connection should not use an HTTP tunnel. You can use this option for NetBackup clients that are not in a DMZ. By default, the <code>WEB_SERVER_TUNNEL_USE</code> option is not present in the configuration file. This option should appear only once in the configuration file.
Equivalent NetBackup web UI property	No equivalent exists.

Configuring server groups

This chapter includes the following topics:

- [About NetBackup server groups](#)
- [Add a server group](#)
- [Delete a server group](#)

About NetBackup server groups

A server group is a group of NetBackup servers that are used for a common purpose.

A NetBackup **Media sharing** group is a server group that shares tape media for write purposes (backups). All members of a **Media sharing** server group must have the same NetBackup primary server.

A **Media sharing** group can contain the following:

- NetBackup primary server
- NetBackup media servers
- NDMP tape servers

See [“About media sharing”](#) on page 525.

See [“Configure media sharing with a server group”](#) on page 526.

Add a server group

A server group is a group of NetBackup servers that are used for a common purpose. Servers can be in more than one group.

Caution: NetBackup allows a server group name to be the same as the name of a media server. To avoid confusion, do not use same name for a server group and a media server.

To add a server group

- 1 On the left, click **Storage > Media servers**.
- 2 Click **Server groups**.
- 3 Click **Add server group**.
- 4 Provide the information for the server group.

Server group name	Provide a unique name for the server group. Do not use the name for an existing media server or other host. You cannot change the name of an existing server group.
Server group type	Select the type of server group.
State	Active. The server group is available for use. Inactive. The server group is not available for use.
Description	Provide a description of the group.

- 5 To add a server to the group, click **Add**, select the server, then click **Add**.
To remove a server from the group, select the server and click **Remove**.
- 6 Click **Save**.

Delete a server group

You can delete a server group if it is no longer in use. Or, if the purpose of the servers in the group has changed.

To delete a server group

- 1 On the left, click **Storage > Media servers**.
- 2 Click **Server groups**.
- 3 Select the group to delete. Then click **Delete > Delete**.

Enabling support for NAT clients and NAT servers in NetBackup

This chapter includes the following topics:

- [About NAT support in NetBackup](#)
- [Important notes](#)
- [Workflow to enable NAT hosts in NetBackup domain](#)
- [Configuring the NetBackup Messaging Broker service](#)
- [Removing NAT support from NetBackup](#)
- [Communication with clients other than NAT clients](#)
- [Performance characteristics of NAT support](#)

About NAT support in NetBackup

NetBackup supports NetBackup clients and servers in a private network that are connected to NetBackup servers in a public network via a device that performs Network Address Translation (NAT). This document refers to such NetBackup clients and servers as NAT clients and NAT servers respectively.

NAT clients and NAT servers together are referred to as NAT hosts.

NetBackup supports NAT clients and NAT servers (or a NAT host) in a network topology where the following conditions are met:

1. A NAT host should be able to resolve the host names of the NetBackup servers that are deployed in a public network and initiate connections with them. It is not required that the NetBackup servers be able to initiate connections to the NAT host.
2. A host name assigned to a NAT host should be resolvable in the private network. It is not required that the host name of the NAT host be resolvable from the NetBackup servers in the public network.
3. Bi-directional connectivity should exist between the primary server and all media servers.
4. Bi-directional connectivity is required between media servers and clients that are behind NAT.
5. The NetBackup software on the NetBackup servers and NAT hosts must be configured for NAT support as described in this document.

When working with NAT hosts, NetBackup software ensures that all network connections are initiated from the NAT client to the NetBackup servers in the public network. In other words, no connections are directly initiated from the NetBackup servers to the NAT hosts. The NAT host support relies on a new NetBackup Messaging Broker (`nbmqbroker`) service on the primary server and a subscriber service on each NAT host that maintains a persistent connection to the messaging broker service on the primary server. This enables the NetBackup servers to send commands to the NAT hosts via the messaging service. When a NetBackup server needs to connect to a NAT client (for example to perform a backup) it sends a 'reverse connection request' message to the NAT host via the primary server. On receiving this message, the NAT client initiates a connection to the requesting NetBackup server.

Here is how a connection between a media server and a NAT client takes place:

1. The NetBackup Messaging Broker (`nbmqbroker`) service starts on the primary server if NAT support is enabled.
2. The subscriber service starts on the NAT host along with other client services and subscribes to `nbmqbroker` service on the primary server if NAT support is enabled on the host.
3. When a media server wants to connect to a NAT client or a primary server wants to connect to a NAT server, it publishes the NAT host's reverse connection request message to the message broker that exists on the primary server.
4. The message broker delivers the message to the subscriber service on the NAT host.

5. The subscriber service initiates a connection from the NAT host to the requesting NetBackup server.
6. The media server uses this connection to communicate with the NAT client or the primary server uses this connection to communicate with the NAT server.

See [“Workflow to enable NAT hosts in NetBackup domain”](#) on page 379.

Support for client-initiated connections in NetBackup

NetBackup NAT support can also be used in the following non-NAT environments where it is desirable or mandatory for the NetBackup clients to initiate all connections to the NetBackup servers:

- Clients or servers are behind a firewall that is configured to disallow incoming connections
- Host names of clients or servers cannot be resolved to an IP address from the NetBackup servers, for example DHCP clients without a Dynamic DNS
- Clients or servers to which media servers or primary servers cannot directly connect for any reason

Important notes

Review the following notes while you enable support for NAT hosts in NetBackup.

- Replication target host should be reachable from the source media server.
- Deduplication from the media server in a public network to a private network is not supported, however the reverse is supported.
- Optimized duplication does not work for communication between a NAT media server and a media server in a public network. This is because the deduplication engine does not support a reverse connection.
However, if a NAT host is used as a replication target host for both media servers, optimized duplication works.
- In case of Windows platform, ensure that the 8dot3 name file setting is enabled for the volume where NetBackup primary server software will be installed. The `fsutil` command is used to enable the required file setting.
Refer to the following article: [Fsutil 8dot3name](#)
- You must provide an authorization token during NetBackup certificate deployment on a NAT host, irrespective of the certificate deployment security level that is set on the primary server. This is required because the primary server cannot resolve the client host name that is part of the certificate deployment request to the NAT device's IP address from which the request appears to be coming.

- Automatic host ID-to-host name mapping is disabled for NAT hosts. A NAT host should be referenced in backup policies and NetBackup commands using the host name that is already mapped to its host ID. The initial hostname mappings are established for a host during NetBackup certificate deployment or external certificate enrolment. If you want a NAT host to use an alternative name for connection, you have to manually map the required host names using the **Security > Host mappings** node.
- In a NetBackup domain that comprises application hosts such as SharePoint, Microsoft Exchange server, or Application Clusters, the application host name or data availability group (DAG) name may be different than the one that is used during NetBackup installation. In some cases the Fully Qualified Domain Name (FQDN) of the host is used during NetBackup installation. Therefore, connection between the NetBackup server and the client (or application host) may fail. To resolve this issue, map both the names of the NetBackup client using the **Security > Host mappings** node.

For more details on the security certificates and certificate deployment levels, refer to the [NetBackup Security and Encryption Guide](#).

Workflow to enable NAT hosts in NetBackup domain

The following table provides the workflow to enable NAT hosts in a NetBackup domain.

Install or upgrade NetBackup	Install NetBackup on the primary server, media server, and client computers or upgrade the existing software to a supported version of NetBackup.
Prepare primary server for NAT support	Do the following: 1 Configure the NetBackup Messaging Broker (or <code>nbmqbroker</code>) service. See “Configuring the NetBackup Messaging Broker service” on page 380. 2 Set the <code>INITIATE_REVERSE_CONNECTION</code> configuration option to <code>TRUE</code> on the primary server using the <code>nbsetconfig</code> command. 3 Restart the primary server services.

Prepare peer primary server for NAT support	This configuration is applicable in an A.I.R. setup. Do the following: ■ On the NAT-enabled primary server host, add the name of the peer primary server in the <code>SERVER</code> field in the configuration file (<code>bp.conf</code> on UNIX or Windows registry) using the <code>nbsetconfig</code> command. ■ Restart the primary server services.
Prepare media server for NAT support	Do the following: 1 Set the <code>INITIATE_REVERSE_CONNECTION</code> configuration option to <code>TRUE</code> on the media server using the <code>nbsetconfig</code> command. 2 Restart the media server services.
Prepare NetBackup client for NAT support	Do the following: 1 Set the <code>ACCEPT_REVERSE_CONNECTION</code> configuration option to <code>TRUE</code> on the client using the <code>nbsetconfig</code> command. 2 Restart the client services. In case of silent installation, you need to set the <code>ACCEPT_REVERSE_CONNECTION</code> option in the answer file only once so that the configuration takes place simultaneously for all clients. For UNIX, ensure that the <code>NBInstallAnswer.conf</code> file is updated with the <code>ACCEPT_REVERSE_CONNECTION</code> option. For Windows, ensure that the <code>silentclient.cmd</code> script is edited with the required information about the <code>ACCEPT_REVERSE_CONNECTION</code> option. For more information on the silent installation, refer to the NetBackup Installation Guide.

Configuring the NetBackup Messaging Broker service

To enable the NetBackup primary server for NAT client and NAT server support, you must configure the NetBackup Messaging Broker (`nbmqbroker`) service on the primary server. The service is required to initiate the connection between NAT hosts and the NetBackup servers.

To configure the service

- ◆ Run the `configureMQ -defaultPorts` command on the primary server.

For more information on the command, refer to the [NetBackup Commands Reference Guide](#).

In a cluster configuration, run the `configureMQ` command only on the active node.

Run the following commands to enable the cluster to monitor the `nbmqbroker` service that you have added to a NetBackup cluster group:

```
configureMQ -enableCluster  
  
configureMQ -defaultPorts
```

Configuring NetBackup Messaging Broker service to use PBX port

If required user can use the PBX port communication feature for the communication between the primary server and the NetBackup client.

Following are the steps required to set up PBX communication:

On primary server:

- Using the following command stop the `mqbroker` service (or stop all services):
 - For Linux: `/usr/opensv/mqbroker/bin/nbmqbroker terminate`
 - For Windows: `C:\Program Files\Veritas\NetBackup\mqbroker\bin\nbmqbroker.exe -stop`
- Configure `mqbroker` using the following respective command:
 - For Linux: `/usr/opensv/mqbroker/bin/install/configureMQ -defaultPorts`
 - For Windows: `C:\Program Files\Veritas\NetBackup\mqbroker\bin\install\configureMQ.bat -defaultPorts`

In the above commands, `-defaultPorts` sets the default port values for external and internal ports as follows:

Default external port: 13781

Default internal ports: 13780, 13779, 13778

To set the port values other than the default values, see the options by running the following command:

For Linux: `/usr/opensv/mqbroker/bin/install/configureMQ --help`

For Windows: `C:\Program`

`Files\Veritas\NetBackup\mqbroker\bin\install\configureMQ --help`

- Add the above configured external port to the configuration file.
 Default: 13781
`NB_MQBROKER_PORT = 13781`
 For Linux, add the entry in `bp.conf` file.
 For Windows, add `DWORD` entry in registry editor.
- Using the following command to start the mqbroker service (or start all services):
 - For Linux: `/usr/opensv/mqbroker/bin/nbmqbroker start`
 - For Windows: `C:\Program Files\Veritas\NetBackup\mqbroker\bin>nbmqbroker.exe -start`

On NetBackup client:

- Stop all the services.
- To use PBX port, add the entry in `bp.conf` file:
`NetBackup_CLIENT_VIA_PBX = 1`
 For Linux, add the entry in `bp.conf` file.
 For Windows, add `DWORD` entry in registry editor.
- Start all the services.

In case you happened to run the configure MQ command twice, the fix the issue using the following steps:

- Delete the `/usr/opensv/var/global/mqbroker/mqhome/.erlang.cookie` file.
- Configure the PBX connection again using the above mentioned steps.

Removing NAT support from NetBackup

Use this section to remove the NAT support from NetBackup.

To remove NAT support

- 1 Ensure that no NAT hosts exist in your NetBackup domain.
- 2 Use the `nbsetconfig` command to set the `INITIATE_REVERSE_CONNECTION` configuration option to `FALSE` on the primary server and the media server.

Communication with clients other than NAT clients

When NAT support is enabled on a NetBackup server, the default behaviour of the server is to perform all client connections via the NetBackup messaging broker

service on the primary server. This requires all clients that the server communicates with to have NAT support enabled. In this default configuration, the server will fail to communicate with any clients that have NAT support disabled or those running a version of NetBackup lacking NAT support. It is possible to instruct the server to attempt a direct connection to such clients by setting the

`ENABLE_DIRECT_CONNECTION` option to `TRUE`. Setting this option allows a NetBackup server to work with clients that have NAT enabled (using reverse connections) and those that have NAT disabled (using direct connections).

In an A.I.R. setup, both source and target primary servers must be enabled for direct connection.

Performance characteristics of NAT support

Since NAT support can be used to backup and restore NetBackup clients across insecure networks like the internet, data channel encryption is enabled by default for communication with NAT clients and servers (or NAT hosts). This follows the 'secure by default' principle.

NetBackup does not currently offer data channel encryption for the hosts for which NAT support is disabled. Data channel encryption secures the data in-flight between the NAT host and the NetBackup server and does not encrypt the data at-rest. The data channel is secured using the secure communications infrastructure.

The current implementation of data channel encryption incurs significant performance overhead. You can disable data channel encryption for NAT hosts that do not communicate with NetBackup servers over an insecure network.

Set the `ENABLE_DATA_CHANNEL_ENCRYPTION` configuration option to `FALSE` on a NAT host to disable data channel encryption.

When data channel encryption is disabled, the backup and restore performance of NAT hosts is similar to the hosts for which NAT support is disabled.

Configuring host credentials

This chapter includes the following topics:

- [About configuring credentials](#)
- [About configuring Snapshot Management server in NetBackup](#)

About configuring credentials

Credentials appears only if a feature that requires external credentials is licensed.

Use **Media and Device Management > Credentials** to manage log on credentials for the following:

- Cloud storage credentials.
Configure the credentials when you configure the storage server.
See the [NetBackup Cloud Administrator's Guide](#).
- NetBackup Deduplication Engine credentials.
Create the credentials when you configure the storage server.
See the [NetBackup Deduplication Guide](#).
- NDMP hosts.
See the [NetBackup for NDMP Administrator's Guide](#).
- OpenStorage storage servers.
Configure the credentials when you configure the storage server.
See the [NetBackup OpenStorage Solutions Guide for Disk](#).
See the [NetBackup Replication Director Solutions Guide](#).
- Virtual machine server credentials.
See the [NetBackup for VMware Administrator's Guide](#).

- WebSocket Server credentials
See the [NetBackup WebSocket Service \(NBWSS\) Reference Guide](#).
- Nutanix Acropolis Cluster credentials

About configuring Snapshot Management server in NetBackup

You can configure the Snapshot Manager server as a snapshot management server. To configure the Snapshot Manager server in NetBackup you need to add the credentials of the Snapshot Manager server.

You can configure the Snapshot Manager server one of the following:

- NetBackup Administration Console
- The `tpconfig` command line. Refer to the *NetBackup Commands Reference Guide*.
- NetBackup web UI. Refer to the *NetBackup™ Web UI Cloud Administrator's Guide*.

Registering a Snapshot Manager server in NetBackup

To register a Snapshot Manager server as snapshot management server in NetBackup you need to add the credentials of the Snapshot Manager server. It is recommended that you add, update, or delete the Snapshot Manager server from NetBackup.

To register the Snapshot Manager server

- 1 Log on to the **NetBackup web UI**.
- 2 In the left navigation pane, go to **Workloads > Cloud > Snapshot Managers**.
The **Snapshot Managers** pane is displayed.
- 3 Click **Add**.
- 4 Provide IP address or FQDN of the Snapshot Manager if DNS server is configured.
- 5 Provide the **Port** details. The default port value is 443.
- 6 Click **Save**.

Note: The host name address must be DNS resolvable.

To register the Snapshot Manager server on-premises

- 1 Log on to the **NetBackup web UI**.
- 2 In the left navigation pane, go to **Storage > Snapshot Managers**.
The **Snapshot Managers** pane is displayed.
- 3 Click **Add**.
- 4 Provide IP address or FQDN of the Snapshot Manager if DNS server is configured.
- 5 Provide the **Port** details. The default port value is 443.
- 6 Click **Save**.

Note: The host name address must be DNS resolvable.

To update the Snapshot Manager server credentials

- 1 Log on to the **NetBackup Administration Console**.
- 2 In the left navigation pane, go to **Media and Device Management > Credentials > Snapshot Management Server**.
The **Snapshot Server Management** pane is displayed.
- 3 In the **Host Name** column, right-click on the server you want to update.
- 4 In the menu, click **Change**.
- 5 In the **Change Snapshot Manager server** dialog box, update the credentials.
- 6 Enter the Snapshot Manager server user name and password.
- 7 Select the cloud providers you want to associate with the Snapshot Manager server.

For on-premises deployment, select None.

Note: You can associate multiple providers with a server, but you cannot associate multiple servers with the same provider.

- 8 Click **OK**.
A success message is displayed.
- 9 Click **OK**

Configuring Snapshot Manager plug-ins in NetBackup

The Snapshot Manager plug-in you have installed on the Snapshot Manager server must be registered and configured in NetBackup with the associated Snapshot Manager server.

To register a Snapshot Manager plug-in

- 1** Log on to the **NetBackup Administration Console**.
- 2** In the left navigation pane, go to **Media and Device Management > Credentials > Snapshot Management Server**.
The **Snapshot Server Management** pane is displayed.
- 3** Click on the server where you want to add the plug-in. The **Snapshot Manager plugins** pane is refreshed.
- 4** In the **Snapshot Manager plugins** pane, right-click on a blank row.
- 5** In the menu, click **New Snapshot Manager plugin**.
- 6** In the **Add New Plugin** dialog box, from the **Available Plugins** list, select the plug-in you want to add.
- 7** Click **Next**.
- 8** In the **Configure CloudPlugin** dialog box, enter the plug-in ID.

Note: The plug-in ID must be unique and must comprise of **A-Z, a-z, 0-9, +, ., -, _** characters.

- 9** Enter the credential details.

Note: The fields are different for different plug-in types. Refer to the *Snapshot Manager Install and Upgrade Guide* for more information about plug-ins and their parameters.

- 10** Click **OK**.
A success message is displayed.
- 11** Click **OK**.
The newly added plug-in is listed in the **Snapshot Manager plugins** pane.

To modify Snapshot Manager plug-in credentials

- 1 Log on to the **NetBackup Administration Console**.
- 2 In the left navigation pane, go to **Media and Device Management > Credentials > Snapshot Management Server**.
The **Snapshot Server Management** pane is displayed.
- 3 Click in the server under which the plug-in is added. The **Snapshot Manager plugins** pane displays all the associated plug-ins.
- 4 Right-click on the Snapshot Manager plug-in you want to modify.
- 5 In the menu, click **Change Plugin**.
- 6 In the **Selected CloudPlugin** dialog box, update the credential details.

Note: You cannot change the plug-in type and plug-in ID.

Warning: If you enter incorrect credentials, the existing snapshot related information available within Snapshot Manager is lost, which can lead to restore failures. This information can be regenerated if you again enter the right credentials.

- 7 Click **OK**.

Managing media servers

This chapter includes the following topics:

- [Activating or deactivating a media server](#)
- [Adding a media server](#)
- [Registering a media server](#)
- [Deleting all devices from a media server](#)
- [Removing a device host from the NetBackup database](#)
- [About decommissioning a media server](#)
- [Previewing references to a media server](#)
- [Decommissioning a media server](#)

Activating or deactivating a media server

When you activate a media server, NetBackup can use it for backup and restore jobs. You can deactivate a media server. A common reason to do so is to perform maintenance. When a media server is deactivated, NetBackup does not send job requests to it.

When you deactivate a media server, the following things occur:

- Current jobs are allowed to complete.
- If the host is part of a shared drive configuration, it does not scan drives.

To activate or deactivate a media server

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Media servers**. Then click the **Media servers** tab.

- 3 Select the media server to activate or deactivate.
- 4 Click **Activate** or **Deactivate**.

Adding a media server

The following table describes an overview of how to add a media server to an existing NetBackup environment.

Note: The NetBackup Enterprise Media Manager service must be active when a media server is added, devices and volumes are configured, and clients are backed up or restored.

Table 6-1 Adding a media server

Step	Procedure	Section
Step 1	On the new media server host, attach the devices and install any software that is required to drive the storage devices.	See the vendor's documentation.
Step 2	On the new media server host, prepare the host's operating system.	See the NetBackup Device Configuration Guide .
Step 3	On the primary server, add the new media server to the Media servers list of the primary server. Also, add the new media server to the Additional servers list of the clients that the new media server backs up. If the new media server is part of a server group, add it to the Additional servers list on all media servers in the group. Note: Ensure that the name you use in NetBackup is the same as the host name in the TCP/IP configuration.	See "Servers properties" on page 159.
Step 4	Install the NetBackup media server software on the new host.	See the NetBackup Installation Guide .
Step 5	On the primary server, configure the robots and drives that are attached to the media server.	See "Configuring robots and tape drives by using the wizard" on page 425.
Step 6	On the primary server, configure the volumes.	See "About adding volumes" on page 472.
Step 7	On the primary server, add storage units to the media server. Always specify the media server as the media server for the storage unit. The Device Configuration Wizard can create storage units when you configure robots and drives. Therefore, if you created storage units already, skip this step.	See "Creating a storage unit" on page 548.

Table 6-1 Adding a media server (*continued*)

Step	Procedure	Section
Step 8	On the primary server, configure the NetBackup policies and schedules to use the storage units that are configured on the media server.	See “About the Policies utility” on page 649.
Step 9	Test the configuration by performing a user backup or a manual backup that uses a schedule that specifies a storage unit on the media server.	See “Performing manual backups” on page 834.

Registering a media server

If the primary server is not running when you install a media server, the media server is not registered. You cannot discover, configure, and manage the devices of that media server. You must register the media server with the primary server.

To register a media server

- 1 Start the EMM service on the primary server.
- 2 On the primary server, run the following command. (For *hostname*, use the host name of the media server.)

On Windows:

```
install_path\NetBackup\bin\admincmd\nbemcmd -addhost -machinename  
hostname -machinetype media -masterserver server_name  
-operatingsystem os_type -netbackupversion  
level.major_level.minor_level
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbemcmd -addhost -machinename  
hostname -machinetype media -masterserver server_name  
-operatingsystem os_type -netbackupversion  
level.major_level.minor_level
```

Note: Ensure that the name you use in NetBackup is the same as the host name in the TCP/IP configuration.

For more information about `nbemcmd` command usage, see the [NetBackup Commands Reference Guide](#).

Deleting all devices from a media server

You can delete all devices from a media server. The media server can be up, down, or failed and unrecoverable. All devices include robots, drives, and disk pools.

Two procedures exist: one to delete all robots and drives and the other to delete disk pools.

To delete all robots and drives from a media server

- ◆ Enter the following command on the primary server:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbemcmd -deletealldevices  
-machinename server_name -machinetype media
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbemcmd -deletealldevices  
-machinename server_name -machinetype media
```

Replace *server_name* with the name of the media server.

To delete disk pools from a media server

- 1 If the media server has disk pools configured, remove the media server from the storage units that use those disk pools. For each storage unit, run the following command on the primary server:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpsturep -label
storage_unit_label -delhost host_name
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpsturep -label
storage_unit_label -delhost host_name
```

Replace *storage_unit_label* with the name of the storage unit and *host_name* with the name of the media server.

- 2 If the media server is the only storage server for the disk pools, change the state of the disk pools to DOWN. To do so, enter the following command on the primary server for each disk pool:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -changestate
-stype server_type -dp disk_pool_name -state DOWN
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -changestate -stype
server_type -dp disk_pool_name -state DOWN
```

Replace *server_type* with the type of storage server: AdvancedDisk, PureDisk, or the vendor string that identifies the OpenStorage server type.

Replace *disk_pool_name* with the name of the disk pool.

- 3 For each disk pool, do the following:

- Remove the media server from disk pool access by entering the following command on the primary server:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -changedp -dp
-disk_pool_name -stype server_type -del_storage_servers
storage_server
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -changedp -dp
disk_pool_name -stype server_type -del_storage_servers
storage_server
```

Replace *disk_pool_name* with the name of the disk pool.

Replace *server_type* with the type of storage server: AdvancedDisk, PureDisk, or the vendor string that identifies the OpenStorage server type. Replace *storage_server* with the name of the media server.

- If the disk pool is on disk storage available only to the media server and is no longer required, delete the disk pool as follows:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -deletedp -dp  
disk_pool_name -stype server_type
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -deletedp -dp  
disk_pool_name -stype server_type
```

You cannot delete a disk pool that has unexpired backup images. You must first expire the images and delete the image fragments, as follows:

- Expire the image as follows:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpexpdate -dp  
disk_pool_name -stype server_type -nodelete
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpexpdate -dp  
disk_pool_name -stype server_type -nodelete
```

- Determine the media IDs in the disk pool as follows:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpimmedia -dp  
disk_pool_name -stype server_type -nodelete
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpimmedia -dp  
disk_pool_name -stype server_type
```

- Delete each media ID in the disk pool as follows:

On Windows:

```
install_path\NetBackup\bin\nbdelete -dt disk_type -media_id  
name
```

On UNIX:

```
/usr/opensv/netbackup/bin/nbdelete -dt disk_type -media_id  
name
```

Removing a device host from the NetBackup database

Use the following procedure to remove a device host from the NetBackup database.

To remove a device host from the NetBackup database

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Media servers**.
- 3 Select the server.
- 4 Select **Delete device host > Delete**.

About decommissioning a media server

You can use the NetBackup `nbdecommission` to decommission a media server. The command launches a text-based wizard that guides you through the decommission process. The wizard removes the references to a media server from a NetBackup domain. (You may have to remove some references manually; the wizard provides instructions to do so in most cases.)

The `nbdecommission` command helps in the following scenarios:

- You retire a server that is no longer required. After all of the backup images on the old server expire, you can use the `nbdecommission` to remove the old server from the NetBackup environment.
- You replace an old server with a new server and keep the same storage. You want to access all of the old server storage and backup images from the new server.
- The old server fails, and you need to replace it with a new server.
- You tried to decommission a media server manually and references to it still remain. The wizard may clean up any references that remain.

Throughout this documentation, the media server to be decommissioned is referred to as the old server.

Warning: Be careful when you use the `nbdecommission` command. Because the command may expire images, data loss may occur. Therefore, you should understand completely what the command does before you use it. It is recommended that you first preview all of the references to a media server before you decommission it.

See [“Previewing references to a media server”](#) on page 401.

The following topics provide more information and provide guidance for the entire process.

See [“About decommissioning limitations”](#) on page 396.

See [“Before you decommission a media server”](#) on page 397.

See [“Decommission actions”](#) on page 398.

See [“Post decommission recommendations”](#) on page 397.

See [“Previewing references to a media server”](#) on page 401.

See [“Decommissioning a media server”](#) on page 402.

About decommissioning limitations

The following are the limitations of the `nbdecommission` command:

- Does not decommission clustered media servers. Those include NetBackup failover media servers or application clusters.
- Does not process the Media Server Deduplication Pools.
Before you decommission a media server that hosts a **Media Server Deduplication Pool**, you must deactivate MSDP on that media server. For more information, see the [NetBackup Deduplication Guide](#).
- Does not update the `vm.conf` files on the NetBackup servers in your environment. Therefore, the old server may remain in the `vm.conf` files on the NetBackup servers.
- Does not update the configuration files on the clients. Therefore, the old server may remain in the server lists on the clients. If you replace an old server with a new server, the new server is not added to the server list of the new client.
- Does not process the NetBackup Vault profiles. If NetBackup Vault profiles exist that refer to the storage units on the old server, update the Vault profiles manually.
- Does not notify you about orphaned resources.
- Does not restart the daemons and services on other servers that the decommissioning affects.
- Requires that you shut down all daemons and services on the old server after it is decommissioned.
- Requires that you reconfigure devices on the new server manually (if required).
- Requires that you know which jobs are running on the old server. You must kill them or let them run to completion before you run the decommission process.

- The `-list_ref` option only reports on the references that it removes explicitly. The command removes some items implicitly and it does not report them. For example, host aliases and host credentials are removed but not reported.
- Requires that you move any media ID generation rules that exist on the old server. You must move them manually to the media server that performs robot inventory.
- Moves the old server to an Administrative Pause state so that no new jobs are started. However, NetBackup still can start backup and restore jobs for basic disk; they obtain resources differently than do jobs for other storage destinations. Also, the `nbdecommission` command may clear the Administrative Pause to expire images (depending on your responses to the wizard). Jobs may start during this period.

Before you decommission a media server

Before you decommission a media server, it is recommended that you do the following:

- Preview the actions of the `nbdecommission` command.
See [“Previewing references to a media server”](#) on page 401.
Analyze the output of the preview operation to ensure that the command captures all references to the old server. If it did not, make a list of the items that the command does not cover and fix them manually later.
- Back up the NetBackup catalog before you begin. You can use it to return your environment to the pre-decommission state if something goes wrong or you have to abort the decommission.
- Run the command during a maintenance window when the load on the NetBackup environment is minimal.

Post decommission recommendations

The following actions are recommended after you run the `nbdecommission` command:

- Follow all of the instructions the command provides.
The command may provide instructions for performing the actions that it cannot perform. For example, it may provide instructions to cancel the backup jobs that are active on the old server.
- Move the physical storage (if needed) and then reconfigure and reinventory those devices.

- Examine the `vm.conf` files on all of the NetBackup servers in your environment. Remove references to the old server and add references to the new server where necessary.
- Remove the old server from the server lists on the clients and add the new server where necessary.
The `nbdecommission` command outputs a list of clients that refer to old server.
- Verify that the old server was removed correctly. Examine the various logical components (backup policies, storage units, and so on) to make sure that the old server references have been removed.
- Back up the NetBackup catalog as soon as possible.
- Uninstall NetBackup from the media server or remove that media server from the environment. After this configuration change, the media server daemons do not have permissions to communicate with the primary server and should no longer be started.

Decommission actions

The `nbdecommission` command deletes the configuration for the old server from the NetBackup database, the NetBackup image catalog, and configuration files on servers.

The following table shows the actions it performs for the components that reference the media server. The table is organized in the order in which the command processes the component.

Table 6-2 `nbdecommission` command actions

Component	Action
Storage unit - Tape	Deletes the following tape storage units: ■ Those in which the Storage device attribute specifies a robot for which the old server is the robot control host. ■ Those in which the Media server attribute specifies the old server. ■ Those in which the Media server attribute specifies Any Available and the old server is the only server that can access the storage unit.
Tape drive	Deletes the tape drive path for each tape drive that is attached to the old server. If the path on the old server is the only path, it also deletes the tape drive. If a path to a drive exists on more than one media server, the tape drive may become unusable. You may have to connect the tape drive to a different media server and then reconfigure it in NetBackup. For example, if the old server is a scan host for a shared drive, NetBackup cannot use the drive if no other host can scan.

Table 6-2 `nbdecommission` command actions *(continued)*

Component	Action
Robotic library	Deletes all of the robotic libraries that are attached to the old server. If the old server is the robot control host for a shared library, the drives and media become standalone and unusable. You must reconfigure and re-inventory the library.
Tape media	Specifies if you want to expire the following tape media or move them to another media server: ■ Those assigned to the old server. ■ Those owned by a media sharing group in which the old server is the only member of the group. ■ Those that have no specific Media owner and the last write host is same as the old server.
Storage unit - BasicDisk	Deletes the storage unit if no images exist on it. If images exist, the wizard lets you choose one of the following options: ■ Expire the images and delete the storage unit. ■ Move the images to the new server. The wizard also updates the Media server field in the storage unit. The BasicDisk storage must be shared, and the same disk path must be available on the new server.
Storage unit - AdvancedDisk	Specifies that if more than one media server can access the disk pool that is the destination of the storage unit, it does the following: ■ Removes the old server from the Media Servers list of the storage unit. ■ Deletes the old server as a storage server. If the old server is the only server that can access the disk pool, the wizard lets you choose to do one of the following: ■ Move the storage and images to the new server and delete the old server as a storage server. The disk volumes must be available on the new server at the same path as the old server. ■ Expire the images (if any), delete any storage units that reference the disk pool, delete the disk pool, and delete the storage server. (A reference is when the disk pool appears in the Select disk pool setting of a storage unit.)

Table 6-2 `nbdecommission` command actions (*continued*)

Component	Action
Storage unit - OpenStorage	Specifies that if more than one media server can access the disk pool that is the destination of the storage unit, it does the following: ■ Removes the old server from the Media Servers list of the storage unit. ■ Deletes the media server as an OpenStorage storage server. If the old server is the only server that can access the disk pool, the wizard lets you choose to do one of the following: ■ Transfer the credentials to the new server and update the Media server field in the storage unit if required. ■ Expire the images (if any), delete any storage units that reference the disk pool, and delete the disk pool. (A reference is when the disk pool appears in the Select disk pool setting of a storage unit.)
Storage unit group	Specifies that if the <code>nbdecommission</code> command deletes all of the storage units in a storage unit group, it also deletes the storage unit group. Deleting the storage unit group also may affect backup policies and storage lifecycle policies. See “Backup policy and schedule” and “Storage lifecycle policy” in this table.
Backup policy and schedule	Deactivates any backup policy in which the storage destination (directly or indirectly) is a storage unit that the command deletes. Specifically, deactivates any backup policy that meets any of the following conditions: ■ The destination is a storage unit that the <code>nbdecommission</code> command deleted. ■ The destination is a storage unit group that contains only one storage unit and the <code>nbdecommission</code> command deleted that storage unit. ■ The destination is a storage lifecycle policy and the <code>nbdecommission</code> command deleted the storage unit that is a Backup operation of the storage lifecycle policy.
Storage lifecycle policy	Specifies that for each storage lifecycle policy in which one or more operations uses a storage unit that the command deleted, it does the following: ■ If images under the SLP control are in-process or yet to be processed, displays the commands to cancel the SLP jobs and then exits. After you cancel the jobs (or wait until the jobs complete), rerun the <code>nbdecommission</code> command to continue with the decommissioning. ■ If all of the images under SLP control are processed, deactivates the storage lifecycle policy. ■ If a deleted storage unit was used by a Backup or Snapshot operation, deactivates all backup policies with the storage lifecycle policy as the destination.
Fibre Transport media server	Displays the commands necessary to delete the old server as an FT media server and then exits. After you delete the old server as an FT media server, rerun the <code>nbdecommission</code> command to continue with the decommissioning.

Table 6-2 nbdecommission command actions *(continued)*

Component	Action
bp.conf file	On UNIX NetBackup servers, removes the old server from the following bp.conf file entries: ■ SERVER ■ MEDIA_SERVER ■ CLIENT_NAME ■ BROWSER On UNIX primary servers, also removes the old server from the FORCE_RESTORE_MEDIA_SERVER and FAILOVER_RESTORE_MEDIA_SERVERS entries.
Windows registry	On Windows NetBackup servers, removes the old server from the following registry keys: ■ SERVER ■ MEDIA_SERVER ■ CLIENT_NAME ■ BROWSER On Windows primary servers, also removes the old server from the FORCE_RESTORE_MEDIA_SERVER and FAILOVER_RESTORE_MEDIA_SERVERS keys.
Clients	Lists the clients on which the old server appears in their server lists. You must remove the references to the old server manually.

Previewing references to a media server

Use the following procedure to preview the associations and references to a media server that you want to decommission. It is recommended that you preview the references to a media server before you decommission it.

The old server does not have to be up and responsive.

See [“About decommissioning a media server”](#) on page 395.

See [“Decommissioning a media server”](#) on page 402.

To preview references to a media server

- 1 Run the `nbdecommission` command on the primary server or on a media server. The following is the command syntax:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdecommission -list_ref  
-oldserver OldServer > file.txt
```

On UNIX:

```
/usr/openv/netbackup/bin/admincmd/nbdecommission -list_ref  
-oldserver OldServer > file.txt
```

Replace *OldServer* with the name of the host to be decommissioned. Replace *file* with a name that denotes its contents or purpose.

- 2 Analyze the output of the preview operation to ensure that the command captures all references to the old server. If it did not, make a list of the items that the command does not cover and fix them manually later.

Decommissioning a media server

This topic is part of a group of topics that provide information about decommissioning a NetBackup media server.

See [“About decommissioning a media server”](#) on page 395.

Your path through the `nbdecommission` wizard depends on how you respond to the wizard prompts. Depending on your environment and how you respond to prompts, the wizard may advise you to perform an action and then exit. To continue in the wizard, you must run the wizard again after you perform the advised action. You may have to exit and rerun the wizard several times.

If active jobs exist on the media server, you must cancel them before the command can begin to decommission the media server. Alternatively, you can wait until they finish.

The *OldServer* does not have to be up and responsive.

It is recommended that you preview the media server references before you decommission a media server.

See [“Previewing references to a media server”](#) on page 401.

The `nbdecommission` command logs to the standard NetBackup administrator commands log directory.

Two procedures exist, as follows:

Replace an old media server with a new media server	See “To replace an old media server with a new media server” on page 403.
Decommission a media server	See “To decommission a media server” on page 404.

To replace an old media server with a new media server

- 1 Run the `nbdecommission` command on the primary server or on a media server that is not the object of this operation. The following is the command syntax:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdecommission -oldserver
OldServer [-newserver NewServer] [-file decom_ops.txt]
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdecommission -oldserver
OldServer [-newserver NewServer] [-file decom_ops.txt]
```

Replace *OldServer* with the name of the host to be decommissioned.

`-newserver` is optional. If you specify a new server, the new server becomes the default media server for the replacement operations. If you do not specify a new server, the wizard prompts you for the new server for each storage type that contains valid backup images. This method is useful if you want to move backup images to different media servers. For example, you can move backup images from tape storage to one media server and backup images from disk storage to another media server.

`-file` is optional. It writes the command operations to the specified file. Replace *decom_ops.txt* with a name that denotes its purpose or contents. It is recommended that you use the `-file` option to maintain a record of the command operations.

- 2 Follow the prompts and perform the requested actions.

For example, the command may make changes on the primary server and on multiple media servers. You may be required to restart the NetBackup services on those servers so that the changes take effect.

To decommission a media server

- 1 Run the following command on the primary server or on a media server that is not the object of this operation. The *OldServer* does not have to be up and responsive.

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdecommission -oldserver  
OldServer
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdecommission -oldserver  
OldServer
```

Replace *OldServer* with the name of the host to be decommissioned.

- 2 Follow the prompts and perform the requested actions.

Configuring storage

- [Chapter 7. Configuring disk storage](#)
- [Chapter 8. Configuring robots and tape drives](#)
- [Chapter 9. Configuring tape media](#)
- [Chapter 10. Inventorying robots](#)
- [Chapter 11. Configuring storage units](#)
- [Chapter 12. Staging backups](#)
- [Chapter 13. Configuring storage unit groups](#)

Configuring disk storage

This chapter includes the following topics:

- [About configuring BasicDisk storage](#)
- [About configuring disk pool storage](#)
- [Configuring NetBackup MSDP disk pools](#)

About configuring BasicDisk storage

A **BasicDisk** type storage unit consists of a directory on locally-attached disk or network-attached disk. The disk storage is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory.

No special configuration is required for **BasicDisk** storage. You specify the directory for the storage when you configure the storage unit.

See [“Creating a storage unit”](#) on page 548.

About configuring disk pool storage

You can configure disk pools if you license a NetBackup feature that uses disk pools.

For more information, see the following guides:

- *The NetBackup AdvancedDisk Storage Solutions Guide.*
- *The NetBackup Cloud Administrator's Guide.*
- *The NetBackup Deduplication Guide.*
- *The NetBackup OpenStorage Solutions Guide for Disk.*
- *The NetBackup Replication Director Solutions Guide.*

Configuring NetBackup MSDP disk pools

You can configure and manage disk pools for NetBackup **Media Server Deduplication Pool** deduplication.

See [“About disk pools for NetBackup deduplication”](#) on page 407.

See [“Configuring a disk pool for deduplication”](#) on page 408.

See [“Managing Media Server Deduplication Pools”](#) on page 410.

About disk pools for NetBackup deduplication

NetBackup deduplication disk pools represent the storage for deduplicated backup data. NetBackup servers or NetBackup clients deduplicate the backup data that is stored in a deduplication disk pool.

Two types of deduplication pools exist, as follows:

- A NetBackup **Media Server Deduplication Pool** represents the disk storage that is attached to a NetBackup media server. NetBackup deduplicates the data and hosts the storage.
NetBackup requires exclusive ownership of the disk resources that comprise the deduplication pool. If you share those resources with other users, NetBackup cannot manage deduplication pool capacity or storage lifecycle policies correctly.

How many deduplication pools you configure depends on your storage requirements. It also depends on whether or not you use optimized duplication or replication, as described in the following table:

Table 7-1 Deduplication pools for duplication or replication

Type	Requirements
Optimized duplication within the same NetBackup domain	Optimized duplication in the same domain requires the following deduplication pools: ■ At least one for the backup storage, which is the source for the duplication operations. The source deduplication pool is in one deduplication node. ■ Another to store the copies of the backup images, which is the target for the duplication operations. The target deduplication pool is in a different deduplication node.

Table 7-1 Deduplication pools for duplication or replication (continued)

Type	Requirements
Auto Image Replication to a different NetBackup domain	Auto Image Replication deduplication pools can be either replication source or replication target. The replication properties denote the purpose of the deduplication pool. The deduplication pools inherit the replication properties from their volumes. See “About the replication topology for Auto Image Replication” on page 967. Auto Image Replication requires the following deduplication pools: ■ At least one replication source deduplication pool in the originating domain. A replication source deduplication pool is one to which you send your backups. The backup images on the source deduplication pool are replicated to a deduplication pool in the remote domain or domains. ■ At least one replication target deduplication pool in a remote domain or domains. A replication target deduplication pool is the target for the duplication operations that run in the originating domain. See “About NetBackup Auto Image Replication” on page 961.

See [“Changing a Media Server Deduplication Pool properties”](#) on page 411.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 414.

Configuring a disk pool for deduplication

Before you can configure a NetBackup disk pool, a NetBackup deduplication storage server must exist.

See [“About disk pools for NetBackup deduplication”](#) on page 407.

When you configure a deduplication disk pool, you specify the following:

- The deduplication storage server to query for the disk storage to use for the pool.
- The disk volume to include in the pool.
NetBackup exposes the storage as a single volume.
- The disk pool properties.

Cohesity recommends that disk pool names be unique across your enterprise.

To configure a deduplication disk pool

- 1 On the left, click **Storage > Disk storage**. Click the **Disk pools** tab, then click **Add**.

Another way to create a disk pool is to click **Create disk pool** at the top of the screen after you have created a storage server.

- 2 In the **Disk pool options**, enter all the required information.

Click **Change** to select a storage server.

If the option **Limit I/O streams** is cleared, the default value is **Unlimited** and may cause performance issues.

Click **Next**.

- 3 In **Volumes**, use the **Volume** drop down to select a volume or add a new volume. If you want to add a new disk pool volume, use the **Add volume** option.

You can configure an MSDP storage server to use cloud storage. Select an existing cloud volume or create a new volume for the MSDP storage server.

Note: When you enable Server-Side Encryption, you can configure AWS Customer-Managed keys. These keys cannot be deleted once they are in use by NetBackup. Each object is encrypted with the key during upload and deleting the key from AWS causes NetBackup restore failures.

Note: FortKnox supports multiple options. For FortKnox Amazon and Amazon Government options and FortKnox Azure and Azure Government. In the web UI, you must contact your Cohesity NetBackup account manager for credentials or with any questions.

For more information on environments and deployment, refer to [FortKnox for NetBackup](#).

For more information about FortKnox Amazon and FortKnox Azure options, refer to the [NetBackup Deduplication Guide](#).

Enter all the required information based on the selection and click **Next**.

- 4** In **Replication**, click **Add** to add replication targets to the disk pool.

This step lets you select a trusted primary server or add a trusted primary server. You can add a primary server that supports NetBackup Certificate Authority (NBCA), ECA, and ECA together with NBCA.

Replication is supported only on MSDP.

Review all the information that is entered for the replication targets and then click **Next**.

- 5** On the **Review** page, verify that all settings and information are correct. Click **Finish**.

The disk pool creation and replication configuration continue in the background if you close the window. If there is an issue with validating the credentials and configuration of the replication, you can use the **Change** option to adjust any settings.

Note: The image sharing server, which already has a disk pool that is created for it, is not available while NetBackup creates a new disk pool for image sharing.

See [“Viewing Media Server Deduplication Pool attributes”](#) on page 413.

Managing Media Server Deduplication Pools

After you configure NetBackup deduplication, you can perform various tasks to manage your deduplication disk pools.

See [“Viewing Media Server Deduplication Pools”](#) on page 411.

See [“Changing a Media Server Deduplication Pool properties”](#) on page 411.

See [“Determining the Media Server Deduplication Pool state”](#) on page 411.

See [“Determining the MSDP disk volume state”](#) on page 411.

See [“Changing the MSDP disk volume state”](#) on page 412.

See [“Viewing Media Server Deduplication Pool attributes”](#) on page 413.

See [“Setting a Media Server Deduplication Pool attribute”](#) on page 414.

See [“Clearing a Media Server Deduplication Pool attribute”](#) on page 415.

See [“Resizing the MSDP storage partition”](#) on page 416.

See [“Deleting a Media Server Deduplication Pool”](#) on page 416.

Viewing Media Server Deduplication Pools

You can view the configured disk pools.

To view disk pools

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Disk storage**.
- 3 Click the **Disk pools** tab.

Changing a Media Server Deduplication Pool properties

You can change the properties of a deduplication disk pool.

To change disk pool properties

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Disk storage**.
- 3 Click the **Disk pools** tab.
- 4 Click the name of the disk pool.
- 5 Click the **Details** tab.
- 6 Click **Edit** and make the wanted changes.
- 7 Click **Save**.

Determining the Media Server Deduplication Pool state

The disk pool state is UP or DOWN.

To determine disk pool state

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Disk storage**.
- 3 Click the **Disk pools** tab.
- 4 Review the **Status** column.

Determining the MSDP disk volume state

Use the NetBackup `nbdevquery` command to determine the state of the volume in a deduplication disk pool. NetBackup exposes all of the storage for MSDP as a single volume, **PureDiskVolume**. The command shows the properties and attributes of the **PureDiskVolume**.

To determine MSDP disk volume state

- ◆ Display the volume state by using the following command:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbdevquery -listdv -stype PureDisk -U -dp disk_pool_name`

Windows: `install_path\NetBackup\bin\admincmd\nbdevquery -listdv -stype PureDisk -U -dp disk_pool_name`

The *state* is either UP or DOWN.

The following is example output

```
Disk Pool Name      : MSDP_Disk_Pool
Disk Type           : PureDisk
Disk Volume Name    : PureDiskVolume
Disk Media ID       : @aaaab
Total Capacity (GB) : 49.98
Free Space (GB)     : 43.66
Use%                : 12
Status              : UP
Flag                : ReadOnWrite
Flag                : AdminUp
Flag                : InternalUp
Num Read Mounts     : 0
Num Write Mounts    : 1
Cur Read Streams   : 0
Cur Write Streams  : 0
```

See [“Changing the MSDP disk volume state”](#) on page 412.

Changing the MSDP disk volume state

The disk volume state is **UP** or **DOWN**. NetBackup exposes all of the storage for MSDP as a single volume, **PureDiskVolume**.

To change the state to **DOWN**, the disk pool in which the volume resides must not be busy. If backup jobs are assigned to the disk pool, the state change fails. Cancel the backup jobs or wait until the jobs complete.

To change the MSDP disk volume state

- ◆ Change the disk volume state; the following is the command syntax:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbdevconfig -changestate
-stype PureDisk -dp disk_pool_name -dv PureDiskVolume -state state`

Windows: `install_path\NetBackup\bin\admincmd\nbdevconfig
-changestate -stype PureDisk -dp disk_pool_name -dv PureDiskVolume
-state state`

For the `-state`, specify either **UP** or **DOWN**.

See [“Determining the MSDP disk volume state”](#) on page 411.

Viewing Media Server Deduplication Pool attributes

Use the NetBackup `nbdevquery` command to view deduplication pool attributes.

To view MSDP pool attributes

- ◆ The following is the command syntax to view the attributes of a deduplication pool. Run the command on the NetBackup primary server or on the deduplication storage server:

UNIX: `/usr/opensv/netbackup/bin/admincmd/nbdevquery -listdp -dp pool_name -stype PureDisk -U`

Windows: `install_path\NetBackup\bin\admincmd\nbdevquery -listdp -dp pool_name -stype PureDisk -U`

The following is example output:

```
Disk Pool Name      : MediaServerDeduplicationPool
Disk Pool Id       : MediaServerDeduplicationPool
Disk Type          : PureDisk
Status             : UP
Flag               : OpenStorage
Flag               : AdminUp
Flag               : InternalUp
Flag               : LifeCycle
Flag               : CapacityMgmt
Flag               : OptimizedImage
Raw Size (GB)      : 235.76
Usable Size (GB)   : 235.76
Num Volumes        : 1
High Watermark     : 98
Low Watermark      : 80
Max IO Streams     : -1
Storage Server     : DedupeServer.example.com (UP)
```

This example output is shortened; more flags may appear in actual output.

Setting a Media Server Deduplication Pool attribute

You may have to set attributes on your existing media server deduplication pools. For example, if you set an attribute on the storage server, you may have to set the same attribute on your existing deduplication disk pools.

To set a MSDP disk pool attribute

- 1 The following is the command syntax to set a deduplication pool attribute. Run the command on the primary server or on the storage server.

`nbdevconfig -changedp -dp pool_name -stype PureDisk -setattribute attribute`

The following describes the options that require the arguments that are specific to your domain:

<code>-changedp</code> <code>pool_name</code>	The name of the disk pool.
<code>-setattribute</code> <code>attribute</code>	The <i>attribute</i> is the name of the argument that represents the new functionality. For example, OptimizedImage specifies that the environment supports the optimized synthetic backup method.

The following is the path to the `nbdevconfig` command:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
- Windows: `install_path\NetBackup\bin\admincmd`

2 To verify, view the disk pool attributes.

See [“Viewing Media Server Deduplication Pool attributes”](#) on page 413.

See [“About disk pools for NetBackup deduplication”](#) on page 407.

Clearing a Media Server Deduplication Pool attribute

You may have to clear attributes on your existing media server deduplication pools.

To clear a Media Server Deduplication Pool attribute

- ◆ The following is the command syntax to clear a deduplication pool attribute. Run the command on the primary server or on the storage server.

```
nbdevconfig -changedp -dp pool_name -stype PureDisk
-clearattribute attribute
```

The following describe the options that require your input:

<code>-changedp</code> <code>pool_name</code>	The name of the disk pool.
<code>-setattribute</code> <code>attribute</code>	The <i>attribute</i> is the name of the argument that represents the new functionality.

The following is the path to the `nbdevconfig` command:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
- Windows: `install_path\NetBackup\bin\admincmd`

Resizing the MSDP storage partition

If the volume that contains the deduplication storage is resized dynamically, restart the NetBackup services on the storage server. You must restart the services so that NetBackup can use the resized partition correctly. If you do not restart the services, NetBackup reports the capacity as full prematurely.

To resize the MSDP storage

- 1 Stop all NetBackup jobs on the storage on which you want to change the disk partition sizes and wait for the jobs to end.
- 2 Deactivate the media server that hosts the storage server.
- 3 Stop the NetBackup services on the storage server.
Be sure to wait for all services to stop.
- 4 Use the operating system or disk manager tools to dynamically increase or decrease the deduplication storage area.
- 5 Restart the NetBackup services.
- 6 Activate the media server that hosts the storage server.
- 7 Restart the deduplication jobs.

Deleting a Media Server Deduplication Pool

You can delete a disk pool if it does not contain valid NetBackup backup images or image fragments. If it does, you must first expire and delete those images or fragments. If expired image fragments remain on disk, you must remove those also.

If you delete a disk pool, NetBackup removes it from your configuration.

If a disk pool is the storage destination of a storage unit, you must first delete the storage unit.

To delete an MSDP disk pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Disk storage**.
- 3 Select a disk pool.
- 4 Click **Delete > Yes**.

Configuring robots and tape drives

This chapter includes the following topics:

- [NetBackup robot types](#)
- [About the device mapping files](#)
- [Downloading the device mapping files](#)
- [About configuring robots and tape drives in NetBackup](#)
- [Configuring robots and tape drives by using the wizard](#)
- [Updating the device configuration by using the wizard](#)
- [Adding a robot to NetBackup manually](#)
- [Managing robots](#)
- [Adding a tape drive to NetBackup manually](#)
- [Configure drive name rules](#)
- [Adding a tape drive path](#)
- [Adding a shared tape drive to a NetBackup environment](#)
- [Correlating tape drives and SCSI addresses on Windows hosts](#)
- [Correlating tape drives and device files on UNIX hosts](#)
- [Managing tape drives](#)
- [Performing device diagnostics](#)

- [Verifying the device configuration](#)
- [About automatic path correction](#)
- [Enabling automatic path correction](#)
- [Replacing a device](#)
- [Updating device firmware](#)
- [About the NetBackup Device Manager](#)
- [About external access to NetBackup controlled devices on UNIX](#)
- [Stopping and restarting the device manager](#)

NetBackup robot types

A robot is a peripheral device that moves tape volumes into and out of tape drives. NetBackup uses robotic control software to communicate with the robot firmware.

NetBackup classifies robots according to one or more of the following characteristics:

- The communication method the robotic control software uses; SCSI and API are the two main methods.
- The physical characteristics of the robot. Library refers to a large robot, in terms of slot capacity or number of drives.
- The media type commonly used by that class of robots. HCART (1/2-inch cartridge tape) is an example of a media type.

The table lists the NetBackup robot types that are supported in release 11.1, with drive and slot limits for each type.

To determine which robot type applies to the model of robot that you use, see the [NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List](#) for your release.

Table 8-1 NetBackup robot types in release 11.1

Robot type	Description	Drive limits	Slot limits	Note
ACS	Automated Cartridge System	1680	No limit	API control. The ACS library software host determines the drive limit.
TLD	Tape library DLT	No limit	32000	SCSI control.

Note: The user interface for NetBackup may show configuration options for the peripheral devices that are not supported in that release. Those devices may be supported in an earlier release, and a NetBackup primary server can manage the hosts that run earlier NetBackup versions. Therefore, the configuration information for such devices must appear in the user interface. The NetBackup documentation may also describe the configuration information for such devices. To determine which versions of NetBackup support which peripheral devices, see the [NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List](#).

About the device mapping files

NetBackup uses several files to determine which protocols and settings to use to communicate with storage devices. NetBackup also uses the files during device discovery and configuration.

The device mapping files are available for download from the following URL:

https://www.veritas.com/support/en_US/article.100040093

The download packages contain the following files:

- `external_robotics.txt`
- `external_types.txt`
- `Readme.txt`

In some cases, you can add support for new or upgraded devices without waiting for a release update. To do so, download the current device mapping files package from the Cohesity Technical Support website and configure NetBackup to use that file. For instructions, see the `Readme.txt` file that is supplied with the device mapping file package.

Note: The contents of the device mapping files do not indicate support for any of the devices, only the ability to recognize and automatically configure them.

See [“Downloading the device mapping files”](#) on page 419.

See [“About configuring robots and tape drives in NetBackup”](#) on page 420.

Downloading the device mapping files

Use the following procedure to download the current device mapping files and update the NetBackup database with their information.

See [“About the device mapping files”](#) on page 419.

To download the current device mapping files

- 1
- Go to the following URL:
- https://www.veritas.com/support/en_US/article.100040093
- 2
- In the *NetBackup Device Mappings Files* row in the table , select the link for your operating system.
- A knowledge base article will appear that contains installation instructions and an archive file of the device mappings.
- 3
- Download the archive file, either a .tar or .zip depending on operating system.
- 4
- Follow the instructions in the `Readme.txt` file in the archive to update the device mappings. The `Readme.txt` file contains instructions for both Windows and UNIX operating systems.

About configuring robots and tape drives in NetBackup

Before you configure robots and tape drives in NetBackup, they must be attached to the computer and recognized by the operating system. The server platforms that NetBackup supports may require operating system configuration changes to allow device discovery.

The [NetBackup Device Configuration Guide](#) provides information about how to configure device drivers for the systems that NetBackup supports.

Configure robots and tape drives in NetBackup as follows:

Device Configuration Wizard

It is recommended to use the **Device Configuration Wizard** to add, configure, and update the following types of devices in NetBackup:

- Robots, including those attached to NDMP hosts
- Tape drives, including those attached to NDMP hosts
- Shared drives (for NetBackup Shared Storage Option configurations only)

See [“Configuring robots and tape drives by using the wizard”](#) on page 425.

The wizard discovers the devices that are attached to the media servers and helps to configure them.

See [“About configuring robots and tape drives in NetBackup”](#) on page 420.

Manually

Alternatively, add robots and drives manually as follows:

- Use menu options in the **NetBackup Administration Console**.
 See [“Adding a robot to NetBackup manually”](#) on page 426.
 See [“Adding a tape drive to NetBackup manually”](#) on page 432.
- Use NetBackup commands.
 For more information, see the [NetBackup Commands Reference Guide](#).

Manual methods do not use device discovery.

To add a robot and drives, first add the robot and then add the drives that are in the robot.

Before configuring robots and drives, read the following topics to understand the process.

See [“About device serialization”](#) on page 422.

See [“About adding devices without discovery”](#) on page 423.

See [“About robot control”](#) on page 423.

See [“About drive name rules”](#) on page 424.

See [“Correlating tape drives and device files on UNIX hosts”](#) on page 443.

See [“Correlating tape drives and SCSI addresses on Windows hosts”](#) on page 441.

About device discovery

Device discovery is an exploratory method that determines which peripheral devices a host can detect. Detection depends on physical attachment (SCSI, Fibre Channel, and so on) and device state (on and responding or off and not responding). Detection also depends on host operating system device-layer configuration.

The goal of device discovery is to provide information to enable fully or partially automatic configuration of peripherals for use with NetBackup. Device discovery provides data that correlates the devices that are interconnected across multiple hosts or multiple host bus adapters on the same host.

To discover devices, NetBackup issues SCSI pass-through commands through operating system device files (on UNIX) or APIs (on Windows). The storage devices must be attached to the computer and recognized by the operating system. A pass-through path to a device must exist.

The operating systems that NetBackup supports may require configuration changes to allow device discovery.

The [NetBackup Device Configuration Guide](#) provides information about how to configure device drivers for the systems that NetBackup supports.

NetBackup can discover the following types of devices:

- SCSI-based robotic libraries
- SCSI-based tape drives
- Native parallel SCSI, Fibre Channel Protocol (FCP) and FC-AL (loop) connections
- SCSI over IP (reported)
- API type robots, such as ACS robots
- NDMP devices that run NDMP version 3 or later

See [“Enabling automatic path correction”](#) on page 457.

About device serialization

Device serialization is a firmware feature that allows device identification and configuration. A unique serial number identifies a device.

NetBackup determines device relationships by comparing serial numbers from multiple sources that refer to the same device. If both a robotic library and a drive fully support serialization, NetBackup can determine the drive's position (or address) in the robotic library.

Most robots and drives support device serialization.

If a device supports serialization, the following actions occur when NetBackup queries the device:

- Each robot and each drive return a unique serial number.
- Each robot also returns the number of drives and the serial number for each of the drives in the robot. NetBackup uses the information to determine the correct drive number for each drive in the robot.

If a device does not support serialization, ask the vendor for a new firmware revision that returns serial numbers. Even with the proper firmware, some devices require the vendor to perform other actions to enable serialization for the device.

If you know that the devices do not support serialization, make sure that you follow the maximum configuration limits that the devices allow. You also must coordinate the drives to their device files or SCSI addresses so you can configure them correctly.

See [“Correlating tape drives and SCSI addresses on Windows hosts”](#) on page 441.

See [“Correlating tape drives and device files on UNIX hosts”](#) on page 443.

The more devices in the configuration that do not support serialization, the greater the chance of configuration problems by using the **Device Configuration Wizard**.

About adding devices without discovery

NetBackup supports some devices that cannot be discovered automatically. NetBackup also supports some devices that require user intervention during the discovery process. To add and configure those devices, in the NetBackup web UI select **Storage > Tape storage** or use the `tpconfig` command.

For the devices that NetBackup cannot discover or that do not have serial numbers, automatic device path correction is limited.

About robot control

When you add a robot to NetBackup manually, you must configure how the robot is controlled.

See [“Robot control \(robot configuration options\)”](#) on page 429.

The following table lists the information that is required to configure the robot control types (local, NDMP, and remote). The information that is required depends on the robot type and the media server type.

Table 8-2 Robot control information

Robot type	Media server type	Robot control	Information required for configuration
ACS	Windows, Solaris SPARC, and Linux (except Linux64)	NDMP	NDMP host name and robot device
ACS	All	Remote	ACSL host
TLD	UNIX	Local	Robotic device file
TLD	Windows	Local	Robot device or SCSI coordinates
TLD	Windows, Solaris SPARC, and Linux (except Linux64)	NDMP	NDMP host name and robot device
TLD	All	Remote	Robot control host

Library sharing example


[Figure 8-1](#) shows library sharing with two servers using two drives in a TLD robot.

The robotic control for the robot is on the host that is named eel. One drive in the robot is connected to eel and the other is connected to the host shark.

Host eel is the robot control host. To configure this robot on host eel, select **Robot is controlled locally by this device host**. To configure this robot on host shark,

select **Robot control is handled by a remote host**. Then, enter eel for the **Robot control host**.

Figure 8-1 Robot control host example



TLD robot (HP EML E-Series)

About drive name rules

The drive name rules define the rules NetBackup uses to name drives.

The default, global drive name rule creates names in the following format:

vendor ID.product ID.index

If you use the default global rule when you add Quantum DLT8000 drives, the drives are named as follows: The first one that you add is named QUANTUM.DLT8000.000, the second one QUANTUM.DLT8000.001, and so on.

You can change the default, global drive name rule.

You also can create drive name rules for specific device hosts (each device host can have its own rule). Host-specific rules override the global rule for the devices that are attached to the specified host.

Only one global rule can exist; it is used for all connected device hosts. The global rule is used for the drive name unless a host-specific rule or local rule is specified.

Drive names are limited to 48 characters.

Use any of the following drive attributes as part of a drive name rule:

- Host name
- Robot number
- Robot type
- Drive position

Drive position information varies depending on the robot type. Drive position information can be ACS coordinates or the robot drive number.
- Drive type

- Serial number
- Vendor ID
- Product ID
- Index

A **Custom Text** field is also available which accepts any of the allowable drive name characters.

See [“Configure drive name rules”](#) on page 437.

Configuring robots and tape drives by using the wizard

It is recommended that you use the **NetBackup Device Configuration Wizard** to configure robots and drives. However, you can add robots and drives manually.

To configure robots and drives by using the wizard

- 1 In the **NetBackup Administration Console**, in the left pane, click **Media and Device Management**.
- 2 In the right pane, click the **Configure Storage Devices** and follow the wizard instructions.

The properties you can configure depend on the robot type, the host type, and the robot control.

Updating the device configuration by using the wizard

It is recommended that you use the Device Configuration Wizard to update the NetBackup device configuration when hardware changes occur.

Update the configuration for all storage device changes. For example, if you add or delete a robot or drive or add a new SCSI adapter in a host, update the configuration.

Do not update the device configuration during backup or restore activity.

To update the device configuration by using the wizard

- 1 In the **NetBackup Administration Console**, select **Media and Device Management > Devices**.
- 2 From the list of wizards in the Details pane, click **Configure Storage Devices** and follow the wizard instructions.

Adding a robot to NetBackup manually

When you add a robot manually, you must specify how the robot is controlled.

See [“NetBackup robot types”](#) on page 418.

See [“About robot control”](#) on page 423.


After you add a robot, you should add the robot's drives.

See [“Adding a tape drive to NetBackup manually”](#) on page 432.

Note: It is recommended that you use the **Device Configuration Wizard** to add and update tape storage devices.

To add a robot using the Actions menu

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **New > Robot**.



New Robot

Media Manager host: cayce.rm.com

Device host: cayce.rm.com

Robot type: TLD - Tape Library DLT Robot number: 0

Robot name: TLD(0)

Robot control

☒ Robot is controlled locally by this device host.

☐ Robot control is handled by a remote host.

☐ Robot control is attached to an NDMP host.

Robot device (click browse to select): port 3, Bus 0, Target 0, LUN 0 (FALCON VTL) **Browse...**

OK Cancel Help

- 3 Specify the properties for the robot.

The properties you can configure depend on the robot type, the host type, and the robot control.

See [“Robot properties and configuration options”](#) on page 428.

4 Click **OK**.

The **Stop/Restart Media Manager Device Daemon** dialog box displays.

5 If you intend to make other changes, click **Cancel**. You can restart the Device Manager or the device daemon after you make the final change.

If the device changes are complete, click **OK** to restart the device daemon.

If you restart the device daemon, any backups, archives, or restores that are in progress may also be stopped.

Robot properties and configuration options

This topic describes the robot properties.

Device host

The host to which the device is attached.

Robot type

The type of robot. To locate the robot type to use for specific vendors and models, See the [NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List](#).

Robot number

The unique, logical identification number for the robotic library. This number identifies the robotic library in displays (for example, TLD (21)) and is also used when you add media for the robot.

Note the following:

- Robot numbers must be unique for all robots on all hosts in the configuration, regardless of the robot type or the host that controls them. For example, if you have two robots, use different robot numbers even if different hosts control them.
- If the robot has its robotic control and drives on different hosts, specify the same robot number in all references to that library. That is, use the same robot number on the hosts with the drives as you do on the host that has the robotic control. A tape library DLT robot is one that allows separate robotic control and drive hosts.

Examples are available in the [NetBackup Device Configuration Guide](#).

Robot control

The robot control properties that you can configure depend on how the robot is controlled.

See [“Robot control \(robot configuration options\)”](#) on page 429.

Robot control (robot configuration options)

The **Robot control** section of the prompt specifies the type of control for the robot. The options that you configure depend on the robot type and the media server type.

Table 8-3 Robot configuration properties

Property	Description
Robot is controlled locally by this device host	Specifies that the host to which the robot is attached controls the robot. You must configure other options (depending on the robot type and device host type).
Robot is handled by a remote host	Specifies that a host other than the device host controls the robot. You must configure other options (based on the selected robot type and device host platform).
Robot control is attached to an NDMP host	Specifies that an NDMP host controls the robot. You must configure other options (depending on the robot type and device host type).
ACSLS host	Specifies the name of the Sun StorageTek ACSLS host; the ACS library software resides on the ACSLS host. On some UNIX server platforms, this host can also be a media server. The ACS library software component can be any of the following: ■ Automated Cartridge System Library Software (ACSLS) Examples are available in the NetBackup Device Configuration Guide. ■ STK Library Station ■ Storagenet 6000 Storage Domain Manager (SN6000). This STK hardware serves as a proxy to another ACS library software component (such as ACSLS). Note: If the device host that has drives under ACS robotic control is a Windows server, STK LibAttach software must also be installed. Obtain the appropriate LibAttach software from STK. For compatibility information, see the <i>NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List</i>. https://www.veritas.com/support/en_US/article.100040093
NDMP host name	Specifies the name of the NDMP host to which the robot is attached.
Robot control host	Specifies the host that controls the robot. The name of the host on which the robot information is defined for TLD robots.

Table 8-3 Robot configuration properties (*continued*)

Property	Description
Robot device	The following applies to a Windows device host only. Specifies the name of the robot device. Select Browse and then select a robot from the prompt list displaying the Devices. If the discovery operation fails to discover a robot, select Add manually. Enter either the Port, Bus, Target, and LUN numbers or the device name in the next prompt. If the browse operation fails for any other reason, you are prompted to specify the required information. Use the Windows management tools to find the Port, Bus, Target, and LUN numbers. If the browse operation does not find attached robots, you are prompted to specify the information.
Robotic device file	UNIX device host only. Specifies the device file that is used for SCSI connections. The device files are located in the <code>/dev</code> directory tree on the device host. To specify the robotic device file, select Browse and then select a robotic device file from the list. A list of device files is displayed. Select the appropriate file and select the Select button. If the browse operation fails to show all of the attached robots, click Add manually. Enter the path of the device file in the Robotic device file field and then select the Select button. If the browse operation does not find attached robots, you are prompted to specify the information. Information about how to add device files is available in the NetBackup Device Configuration Guide.
Robot device path	NDMP host only. Specifies the name of the robotic device that is attached to the NDMP host.
Port, Bus, Target, LUN	Windows hosts only. The Port, Bus, Target, and LUN are the SCSI coordinates for the robotic device. To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN.

Managing robots

You can perform various tasks to manage your robots.

See [“Changing robot properties”](#) on page 430.

See [“Delete a robot”](#) on page 431.

Changing robot properties

Use the following procedure to change the configuration information for a robot.

To change the robot control properties of a robot

- 1** Open the **NetBackup web UI**.
- 2** On the left, click **Storage > Tape storage**. Then click the **Robots** tab.
- 3** Select a robot and click **Edit**.
- 4** Change the robot control properties as necessary.

The properties that you can change depend on the robot type, the host type, and the robot control selection.

See [“Changing robot properties”](#) on page 430.
- 5** Click **Save**.

If you restart the Device Manager or the device daemon, any backups, archives, or restores that are in progress also may be stopped.

Delete a robot

Use the following procedure to delete a robot or robots when the media server is up and running.

Any drives that are configured as residing in a robot that you delete are changed to standalone drives.

Any media in the deleted robot is also moved to standalone. If the media is no longer usable or valid, delete it from the NetBackup configuration.

See [“Delete a volume”](#) on page 506.

To delete a robot

- 1** Open the NetBackup web UI.
- 2** On the left, click **Storage > Tape storage**. Click the **Robots** tab.
- 3** Select the robot or robots you want to delete.
- 4** Click **Delete > Delete**.

Note: It may take a few minutes for the web UI to reflect that the robot is deleted. You are prompted to restart the Media Manager device daemon.

Moving a robot and its media to a new media server

Use the following process to move a robot and its media from one server (the *old_server*) to a different media server (the *new_server*).

Table 8-4 Move a robot and media to a new server overview

Task	Procedure
Determine which tapes on the <i>old_server</i> contain NetBackup images that have not expired.	Run the following <code>bpmedialist</code> command: <code>bpmedialist -mlist -l -h old_server</code> The <code>-l</code> option produces one line of output per tape.
Move the tapes in the robot that is attached to the <i>old_server</i> to non-robotic status (standalone).	See “Moving volumes by using the Actions menu” on page 518.
Move the media logically from the <i>old_server</i> to the <i>new_server</i> .	If both the <i>old_server</i> and the <i>new_server</i> are at NetBackup 6.0 or later, run the following command: <code>bpmedia -movedb -allvolumes -oldserver old_server -newserver new_server</code> If either server runs a NetBackup version earlier than 6.0, run the following command for each volume that has active images: <code>bpmedia -movedb -ev media_ID -oldserver old_server -newserver new_server</code> For the media that has active images, see the <code>bpmedialist</code> command output from the first step of this process.
Configure NetBackup so that restore requests are directed to the <i>new_server</i> .	See “Forcing restores to use a specific server” on page 108.
Shut down both the <i>old_server</i> and the <i>new_server</i> .	See the vendor's documentation.
Disconnect the robot from the <i>old_server</i> .	See the vendor's documentation.
Connect the robot to the <i>new_server</i> . Verify that the operating system on the new media server recognizes the robots.	See the vendor's documentation.
Create the appropriate NetBackup storage units.	See “Creating a storage unit” on page 548.
Inventory the robots that are attached to the <i>new_server</i> . The inventory updates the location of all tapes in the robot.	See “Updating the NetBackup volume configuration with a robot's contents” on page 538.

Adding a tape drive to NetBackup manually

Use the following procedures to add a tape drive manually.

Note: It is recommended that you use the **Device Configuration Wizard** to add and update tape storage devices.

To add a drive using the Actions menu

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **New > Tape Drive**.

The screenshot shows the 'New Drive' configuration window. It has a title bar with a gear icon and the text 'New Drive'. The window is divided into several sections:

- Drive name:** A text input field with a 'Configure...' button to its right. Below the input field is a checkbox labeled 'Use drive name rules'.
- Host and path information:** A table with columns: Host, NDMP Host, Path, Port, Bus, Target, LUN, and Enabled. To the right of the table are three buttons: 'Add...', 'Change...', and 'Remove'.
- Drive information:**
 - Drive type:** A dropdown menu currently showing '1/2" Cartridge 2 (hcart2)'.
 - Serial Number:** A text input field.
 - Cleaning frequency (in hours):** A numeric input field with a value of '0' and up/down arrow buttons.
 - ☐ Drive is in a robotic library.
 - Robotic library:** A dropdown menu.
 - Robot drive number:** A numeric input field with a value of '1' and up/down arrow buttons.

At the bottom right of the window are three buttons: 'OK', 'Cancel', and 'Help'.

- 3 For the drive name, do one of the following:
 - Enter a name for the drive in the **Drive name** field.
 See [“Drive name \(tape drive configuration option\)”](#) on page 434.
 - Select **Use drive name rules**. This option uses rules to name the drive automatically.
 See [“About drive name rules”](#) on page 424.
 See [“Configure drive name rules”](#) on page 437.

- 4 To configure the host and the path information, click **Add** in the **Host and path information** area of the dialog box.
 See [“Host and path information \(tape drive configuration options\)”](#) on page 435.
- 5 In the **Drive information** area of the dialog box, configure the drive properties.
 The properties depend on the drive type and host server type.
 See [“Drive information \(tape drive configuration options\)”](#) on page 435.
- 6 After you configure all of the properties, click **OK**.
- 7 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box or the **Media and Device Management** dialog box to restart the Device Manager or the device daemon.
 If you intend to make other changes, click **No**; you can restart the Device Manager or the device daemon after you make the final change.
 If you restart the Device Manager or the device daemon, any backups, archives, or restores that are in progress also may be stopped.

Tape drive configuration options

You can specify properties when you add a tape drive or change the properties of a drive. The properties that you can specify depend on the drive type, server platforms, or NetBackup server types.

Drive name (tape drive configuration option)

Specifies the name of the drive. Each drive name must be unique. It is recommended that you use descriptive names. Drive names are limited to 48 characters.

Alternatively, use the drive name rules to create a unique drive name.

Use drive name rules (tape drive configuration option)

Adds a drive only. Select to use drive name rules to assign names to drives automatically.

To configure drive name rules, click **Configure**.

See [“About drive name rules”](#) on page 424.

See [“Configure drive name rules”](#) on page 437.

Host and path information (tape drive configuration options)

Use the **Host and path information** group box to add or change paths to the drive. You can specify multiple paths to the same physical device. If you specify multiple paths for a drive, it becomes a shared drive.

To add a drive path, click **Add**.

To change a drive path, click **Change**.

To delete a drive path, click **Remove**.

See [“About SCSI reserve on drive paths”](#) on page 439.

See [“Drive path options”](#) on page 439.

Drive information (tape drive configuration options)

The **Drive information** group box includes drive properties. The properties that you can specify depend on the drive type, server platforms, and NetBackup server types.

The following table describes the tape drive configuration options.

Table 8-5 Tape drive configuration options

Option	Description
Drive type	Specifies the type of drive. The following are the valid drive types: ■ DLT (DLT cartridge) ■ DLT2 (DLT cartridge 2) ■ DLT3 (DLT cartridge 3) ■ HCART (1/2-inch cartridge) ■ HCART2 (1/2-inch cartridge 2) ■ HCART3 (1/2-inch cartridge 3)
Drive is in a robotic library	Specifies that the drive is in a robot. If the drive is a standalone drive (it is not in a robot), do not select this option. If you select this option, configure the Robotic library and Robot drive number fields.

Table 8-5 Tape drive configuration options (*continued*)

Option	Description
Cleaning Frequency	Specifies the frequency-based cleaning for the drive. NetBackup does not support drive cleaning in some robot types. If you want to configure a frequency-based cleaning schedule for the drive, set the number of mount hours between each drive cleaning. When you add a drive or reset the mount time to zero, NetBackup records the amount of time that volumes have been mounted in that drive. The default frequency is zero. When the accumulated mount time exceeds the time you specify for the cleaning frequency, drive cleaning occurs if the following are true: ■ If the drive is in a robotic library that supports drive cleaning ■ If a cleaning cartridge is defined in that robotic library ■ If the cleaning cartridge is compatible with the drive that needs to be cleaned ■ If the cleaning cartridge has a nonzero number of cleanings that remain NetBackup resets the mount time when the drive is cleaned. Drives can also be cleaned from the Device Monitor. If you do not specify a cleaning frequency, you can still use automated drive cleaning with the TapeAlert feature. For more information about TapeAlert drive cleaning, see the NetBackup Administrator's Guide, Volume II.
Drive Status	On UNIX only. Specifies the availability of the drive. When you add a drive, the default drive status is UP, which means the drive is available. When a drive is UP, the default mode is AVR (Automatic Volume Recognition). To change the drive status, click UP or DOWN. You can also change the drive status by using the commands on the Actions menu in Device Monitor.
Serial Number	A read-only field that shows the serial number of the drive.
Robotic library	Specifies a robot that controls the drive. You can select any configured robot that can control the drive.

Table 8-5 Tape drive configuration options (*continued*)

Option	Description
Robot drive number	Specifies the physical location in the robot of the drive. When you add more than one drive to a robot, you can add the physical drives in any order. For example, you can add drive 2 before drive 1. The correct robot drive number is critical to the proper mounting and utilization of media. You must determine which logical device name (Windows) or the device file (UNIX) identifies which physical drive in the robot. You should correlate the drive serial number with drive serial number information from the robot, as follows: ■ On Windows: You must determine which physical drive in the robot is identified by the logical device name. See “Correlating tape drives and SCSI addresses on Windows hosts” on page 441. ■ On UNIX: You must determine which physical drive in the robot is identified by the device file name. See “Correlating tape drives and device files on UNIX hosts” on page 443. NetBackup does not detect incorrect drive number assignment during configuration; however, an error occurs when NetBackup tries to mount media on the drive. Note: The Robot drive number property does not apply when you add drives to API robots. API robots are ACS type in NetBackup.
ACS, LSM, Panel, Drive	Specify the drive locations within an ACS robot. The following information applies only to the ACS robot drive. The ACS property specifies the physical location of the drive within the robot. During installation, the correlation between the physical drive in the robot and the device file you specified earlier represents. You establish this correlation during installation. The drive location properties are as follows: ■ ACS Number - specifies the index (in ACS library software terms) that identifies the robot that has this drive. ■ LSM Number - specifies the Library Storage Module that has this drive. ■ Panel Number - specifies the robot panel where this drive is located. ■ Drive Number - specifies the physical number of the drive (in ACS library software terms).

Configure drive name rules

Use the following procedure to configure the rules that NetBackup uses to name tape drives. These rules are defined in the **Configure drives and robots** wizard.

To configure drive name rules

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**.
- 3 At the top right, select **Actions > Configure storage devices**.
- 4 On the **Device hosts** step of the wizard, go to the section **Configured drive name rules**.
- 5 To create a new rule, select **Add**.
- 6 You can replace the global rule or create a local rule:
 - To replace the global rule, select the **Global rule** check box.
A global rule must always exist and only one global rule can exist. If you want to override the existing global rule, you can create another global rule to replace it.
 - To create a local rule, select **Add** and enter or select the name of the device host.
- 7 Select **Next**.
- 8 Configure the rules for naming drives:
 - To select a field that you want to use in the drive name, select **Add field**. Then select the field from the dropdown list.
If you use **host name** in the rule and the drive is a shared drive, the name of the first host that discovers the drive is used as the host name. The name for a shared drive must be identical on all servers that share the drive.
 - To add custom text to the drive name rule, select **Add custom text**. Then enter the custom text.
 - To change the order of the fields that are used for the drive name, select **Actions > Move up** or **Actions > Move down**.
- 9 Select **Create** to finalize the rule.

Adding a tape drive path

Usually, you add a tape drive path when you add a drive to NetBackup. Use the following procedure to add a drive path.

To add a tape drive path

1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Drives**.

2 Select **Add a New Drive**. In the dialog box, click **Add**.

3 In the **Add Path** dialog box, configure the properties for the drive path.

The properties that you can specify depend on drive type, server platform, or NetBackup server type.

See [“About SCSI reserve on drive paths”](#) on page 439.

See [“Drive path options”](#) on page 439.

About SCSI reserve on drive paths

NetBackup lets you configure exclusive access protection to tape drives so that other host bus adaptors (HBAs) cannot control the drives during the reservation. The **Enable SCSI Reserve** host property configures the protection for each media server.

See [“Media properties”](#) on page 119.

For more information about how NetBackup reserves drives, see the [NetBackup Administrator's Guide, Volume II](#).

Drive path options

The following table describes the options to add a drive path.

Table 8-6 Add drive path options

Option	Description
Host name (Windows) Device host (UNIX)	Specifies the device host for the drive.
Enable host path	Specifies that the path is active and that NetBackup can use it for backups and restores.
NDMP host	Specifies the NDMP host for the device (if an NDMP host is configured in your NetBackup environment). For additional information about NDMP drives, see the NetBackup for NDMP Administrator's Guide .

Table 8-6 Add drive path options (*continued*)

Option	Description
Override SCSI Reserve settings	Specifies the SCSI reserve override setting for the drive path. ■ Server Default. Use the SCSI reserve protection setting configured for the media server. If the media server default is no protection, other HBAs can send the commands that can cause a loss of data to the tape drives. ■ SPC-2 SCSI Reserve. This option provides SCSI reserve and release protection for the SCSI devices that conform to the reserve and the release management method. That method is defined in the SCSI Primary Commands - 2 (SPC-2) standard. ■ SCSI Persistent Reserve. This option provides SCSI persistent reserve in and persistent reserve out protection for the SCSI devices that conform to the SCSI Primary Commands - 3 (SPC-3) standard. Global SCSI reserve properties are configured in the Media host properties. See “Media properties” on page 119.
Path	On UNIX hosts. Specifies the path of the character-mode, no rewind device file on the specified host. You can either type-in or browse and select one of the existing devices on the host. The Browse button is not available if the This path is for a Network Attached Storage Device option is selected. See “About no rewind device files on UNIX” on page 440. Device files are in the <code>/dev</code> directory on the UNIX host. If the entries do not exist, see the NetBackup Device Configuration Guide for information about how to create them.
Port, Bus, Target, and LUN	On Windows hosts. You can browse and select one of the existing devices on the host. When you select a device, Port, Bus, Target, and LUN are auto-filled. The Browse button is not available if the This path is for a Network Attached Storage Device option is selected. To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN. The device attributes on Windows systems cannot change during a NetBackup operation.
This path is for a Network Attached storage device	Specifies that the path is for a Network Attached Storage (NAS) device.

About no rewind device files on UNIX

Although both rewind and no rewind on close device files are usually available, NetBackup requires only the no rewind device file. A no rewind device remains at

its current position on a close operation. On some versions of UNIX, the device file name may be preceded or followed by the letter n.

Device files are in the `/dev` directory on the UNIX host. If the entries do not exist, you must create them.

For more information, see the [NetBackup Device Configuration Guide](#).

Adding a shared tape drive to a NetBackup environment

It is recommended that you use the **Configure drives and robots** wizard to add, configure, and update shared drives. This wizard provides the easiest method for adding shared drives in a Shared Storage Option configuration.

For more information about the Shared Storage Option, see the [NetBackup Administrator's Guide, Volume II](#).

Correlating tape drives and SCSI addresses on Windows hosts

If your tape drives do not support device serialization, you may have to determine which logical device name or SCSI address matches the physical drive. You also may have to do so if you add the tape drives manually.

To correlate tape drives and SCSI addresses on Windows hosts

- 1 Note the SCSI target of the drive.
- 2 Correlate the SCSI target to the drive address by using the robot's interface panel. Alternatively, examine the indicators on the rear panel of the tape drive.
- 3 Determine the physical drive address (for example, number) by checking labels on the robot.
- 4 Configure the robot in NetBackup and then add the drives.

When you add the drives, ensure that you assign the correct drive address to each set of SCSI coordinates.

Optionally, use the appropriate NetBackup robotic test utility to verify the configuration.

For more information about the robotic test utilities, see the [NetBackup Troubleshooting Guide](#).

To verify the device correlation Windows

- 1 Stop the NetBackup Device Manager (`ltid`).
- 2 Restart `ltid`, which starts the Automatic Volume Recognition process (`avrd`). Stop and restart `ltid` to ensure that the current device configuration has been activated.

If robotic control is not local to this host, also start the remote robotic control daemon.
- 3 Use the robotic test utility to mount a tape on a drive.
- 4 Use the NetBackup Device Monitor to verify that the tape was mounted on the correct robot drive.

For Windows hosts only.

Assume that a TLD robot includes three drives at the following SCSI addresses:

Drive 1	5,0,0,0
Drive 2	5,0,1,0
Drive 3	5,0,2,0

Windows device correlation example

Also assume that you requested that the tape be mounted on drive 1.

If the SCSI coordinates for the drive are configured correctly, the Administration Console Device Monitor shows that the tape is mounted on drive 1.

If the Device Monitor shows that the tape is mounted on a different drive, the SCSI coordinates for that drive are not correctly configured. For example, if the Device Monitor shows that the tape is mounted on drive 2, the SCSI coordinates for drive 1 are incorrect. Replace the drive 1 SCSI coordinates (5,0,0,0) with the correct SCSI coordinates (5,0,1,0) for drive 2. You also know that the SCSI coordinates for drive 2 are incorrect. Possibly, the SCSI coordinates were swapped during configuration.

Use the robotic test utility to unload and unmount the tape from drive 1. Repeat the test for each drive.

If the data path to the drive in which the tape is mounted is not on the robot control host, you may have to unload the drive. To do so, use a command on another host or use the drive's front panel.

Correlating tape drives and device files on UNIX hosts

If your tape drives do not support device serialization, you may have to determine which device file or SCSI address matches the physical drive. You also may have to do so if you add the tape drives manually.

Correlate device files to physical drives when you create the device files for each drive.

To correlate tape drives and device files on UNIX

- 1 Determine the physical location of each drive within the robotic library. The location usually is shown on the connectors to the drives or in the vendor's documentation.
- 2 Physically connect the drives to SCSI adapters in the host.
- 3 Record the adapter and SCSI addresses to which you connected each drive.
- 4 Create device files for each drive by using the SCSI addresses of the drives and adapters.

Add the device file by using the notes from a previous step to complete the correlation between device files and physical drive location.

- 5 Configure the robot in NetBackup and then add the drives.

When you add the drives, verify that you assign the correct drive address (for example, robot drive number) to each device path.

Optionally, use the appropriate NetBackup robotic test utility to verify the configuration.

For more information about the robotic test utilities, see the [NetBackup Troubleshooting Guide](#).

To verify the device correlation on UNIX

- 1 Stop the NetBackup device daemon (`ltid`).
- 2 Start `ltid`, which starts the Automatic Volume Recognition daemon (`avrd`). Stop and restart `ltid` to ensure that the current device configuration is activated.

If robotic control is not local to this host, also start the remote robotic control daemon.
- 3 Use the robotic test utility to mount a tape on a drive.
- 4 Use the **NetBackup Administration Console Device Monitor** to verify that the tape was mounted on the correct robot drive.

UNIX device correlation example

On UNIX only.

Assume that a TLD robot includes three drives and the operating system includes the following device paths:

Drive 1 /dev/rmt/0cbn

Drive 2 /dev/rmt/1cbn

Drive 3 /dev/rmt/3cbn

Also assume that you requested that the tape be mounted on drive 1.

If the device path for the drive is configured correctly, the **Device monitor** shows that the tape is mounted on drive 1.

If the Device monitor shows that the tape is mounted on a different drive, the device path for that drive is not configured correctly. For example, if the Device monitor shows that the tape is mounted on Drive 2, the device path for drive 1 is incorrect. Replace the drive 1 device path (/dev/rmt/0cbn) with the correct device path (/dev/rmt/1cbn) for drive 2. You may need to use a temporary device path while you make these changes. You also know that the device path for drive 2 is incorrect. Possibly, the device paths were swapped during configuration.

Use the robotic test utility to unload and unmount the tape from drive 1. Repeat the test for each drive.

Additional action is required if the path to the drive where the tape is mounted is not on the host with direct robotic control. In that case you may have to unload the drive with a command from another host or from the drive's front panel.

Managing tape drives

You can perform various tasks to manage tape drives.

To manage tape drives, open the NetBackup web UI. Then on the left click **Storage > Tape storage**.

See [“Adding a tape drive to NetBackup manually”](#) on page 432.

See [“Changing a drive comment”](#) on page 445.

See [“About downed drives”](#) on page 445.

See [“Changing a drive operating mode”](#) on page 446.

See [“Changing a tape drive path”](#) on page 446.

See [“Changing the operating mode for a drive path”](#) on page 448.

See [“Changing tape drive properties”](#) on page 449.

See [“Changing a tape drive to a shared drive”](#) on page 449.

See [“Cleaning a tape drive from the Device monitor”](#) on page 450.

See [“Deleting a drive”](#) on page 450.

See [“Resetting a drive”](#) on page 451.

See [“Resetting the mount time of a drive”](#) on page 451.

See [“Running a tape drive diagnostic test”](#) on page 454.

See [“Setting the drive cleaning frequency”](#) on page 452.

See [“Viewing drive details”](#) on page 452.

Changing a drive comment

You can change the comment that is associated with a drive.

To change a drive comment

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Device monitor** tab.
- 3 Select a drive.
- 4 Select **Actions > Change drive comment**.
- 5 Add a comment or change the current drive comment.
See [“NetBackup naming conventions”](#) on page 1054.
- 6 Click **Save**.

About downed drives

NetBackup downs a drive automatically when there are read or write errors that surpass the threshold within the time window. The default drive error threshold is 2. That is, NetBackup downs a drive on the third drive error in the default time window (12 hours).

Common reasons for write failures are dirty write heads or old media. The reason for the action is logged in the NetBackup error catalog (view the Media Logs report or the All Log Entries report). If NetBackup downs a device, it is logged in the system log.

You can use the NetBackup `nbbemmcmd` command with the `-drive_error_threshold` and `-time_window` options to change the default values.

For more information about `nbbemmcmd`, see the [NetBackup Commands Reference Guide](#).

See [“Changing a drive operating mode”](#) on page 446.

Changing a drive operating mode

Usually you do not need to change the operating mode of a drive. When you add a drive, NetBackup sets the drive state to UP in Automatic Volume Recognition (AVR) mode. Other operating mode settings are used for special purposes.

The drive operating mode is displayed and changed on the **Device monitor** tab.

To change the mode of a drive

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Select a drive or multiple drives.
- 4 Select the operating mode that you want to apply to the drive. For example, to change a drive's status to **Down**, select **Down drive**.

Note that **Up Drive, Operator control** applies only to standalone drives.

- 5 If the drive is configured with multiple device paths or is a shared drive (Shared Storage Option), a screen displays that contains a list of all the device paths to the drive. Select the path or paths to change.

Changing a tape drive path

Use the following procedure to change a drive path.

See [“Changing the operating mode for a drive path”](#) on page 448.

To change a drive path

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Drives**. Double-click the drive that you want to change.
- 2 In the **Change Tape Drive** dialog box, select the drive path in the **Host and Path information** list. Click **Change**.

- 3 In the **Change Path** dialog box, configure the properties for the drive path.
The properties you can change depend on drive type, server platform, or NetBackup server type.
See [“About SCSI reserve on drive paths”](#) on page 439.
See [“Drive path options”](#) on page 439.
- 4 Click **OK** to save the changes.

Drive path options

The following table describes the options to add a drive path.

Table 8-7 Add drive path options

Option	Description
Host name (Windows) Device host (UNIX)	Specifies the device host for the drive.
Enable host path	Specifies that the path is active and that NetBackup can use it for backups and restores.
NDMP host	Specifies the NDMP host for the device (if an NDMP host is configured in your NetBackup environment). For additional information about NDMP drives, see the NetBackup NAS Administrator's Guide .
Override SCSI Reserve settings	Specifies the SCSI reserve override setting for the drive path. ■ Server Default. Use the SCSI reserve protection setting configured for the media server. If the media server default is no protection, other HBAs can send the commands that can cause a loss of data to the tape drives. ■ SPC-2 SCSI Reserve. This option provides SCSI reserve and release protection for the SCSI devices that conform to the reserve and the release management method. That method is defined in the SCSI Primary Commands - 2 (SPC-2) standard. ■ SCSI Persistent Reserve. This option provides SCSI persistent reserve in and persistent reserve out protection for the SCSI devices that conform to the SCSI Primary Commands - 3 (SPC-3) standard. Global SCSI reserve properties are configured in the Media host properties. See “Media properties” on page 119.

Table 8-7 Add drive path options (*continued*)

Option	Description
Path	On UNIX hosts. Specifies the path of the character-mode, no rewind device file on the specified host. You can either type-in or browse and select one of the existing devices on the host. The Browse button is not available if the This path is for a Network Attached Storage Device option is selected. See “About no rewind device files on UNIX” on page 440. Device files are in the <code>/dev</code> directory on the UNIX host. If the entries do not exist, see the NetBackup Device Configuration Guide for information about how to create them.
Port, Bus, Target, and LUN	On Windows hosts. You can browse and select one of the existing devices on the host. When you select a device, Port, Bus, Target, and LUN are auto-filled. The Browse button is not available if the This path is for a Network Attached Storage Device option is selected. To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN. The device attributes on Windows systems cannot change during a NetBackup operation.
This path is for a Network Attached storage device	Specifies that the path is for a Network Attached Storage (NAS) device.

Changing the operating mode for a drive path

The Device monitor shows path information for drives, including the following:

- Multiple (redundant) paths to a drive are configured
- Any drives are configured as shared drives (Shared Storage Option)

To change the operating mode for a drive path

- 1 Open the web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Click on the drive name to view the drive properties. Then click on the **Paths** tab.
- 4 Select a path or select multiple paths.
- 5 Click **Actions**, then choose a command for the path action, as follows:
 - **Up path**
 - **Down path**

- **Reset path**

Changing tape drive properties

Use the following procedure to change the configuration information for a drive.

To change drive properties

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Drives**.
- 2 In the details pane, select the drive you want to change.
- 3 Click **Edit > Change**.
- 4 Change the properties of the drive.
The properties depend on the drive type and host server type.
See [“Tape drive configuration options”](#) on page 434.
- 5 If the device changes are complete, select **Yes** to restart the Device Manager or the device daemon.
If you intend to make other changes, click **No**; you can restart the Device Manager or the device daemon after you make the final change.
If you restart the Device Manager or the device daemon, any backups, archives, or restores that are in progress may also be stopped.
The initial drive status is UP, so the drive is available as soon as you restart the device daemon.
- 6 After you change the properties, click **OK**.

Changing a tape drive to a shared drive

Change a drive to a shared drive by adding paths to a currently configured drive.

To configure and use a shared drive, a Shared Storage Option license is required on each primary server and media server.

To change a drive to a shared drive

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 Select **Drives** in the tree pane.
- 3 Select the drive you want to change in the **Drives** pane.
- 4 Click **Edit > Change**.

- 5 Click **Add**.
- 6 Configure the properties for the hosts and paths that share the drive.

Cleaning a tape drive from the Device monitor

When you add a drive to NetBackup, you can configure the automatic, frequency-based cleaning interval.

You can also perform an operator-initiated cleaning of a drive regardless of the cleaning frequency or accumulated mount time of the drive. However, appropriate cleaning media must be added to NetBackup.

After you clean a drive, reset the mount time.

See [“Resetting the mount time of a drive”](#) on page 451.

Drive cleaning functions can also be performed from the **Activity Monitor**.

See [“Cleaning tape drives from the Activity Monitor”](#) on page 1031.

To clean a tape drive

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Select the drive to clean.
- 4 Select **Actions > Drive cleaning > Clean now**. NetBackup initiates drive cleaning regardless of the cleaning frequency or accumulated mount time.

The **Clean now** option resets the mount time to zero, but the cleaning frequency value remains the same. If the drive is a standalone drive and it contains a cleaning tape, NetBackup issues a mount request.

- 5 For a shared drive (Shared Storage Option), do the following:

In the list of hosts that share the drive, choose only one host on which the function applies.

- 6 Select **Clean now**.

The **Clean now** function can take several minutes to complete, so the cleaning information may not update immediately.

Deleting a drive

Use the following procedure to delete a drive or drives when the media server is up and running.

If the media server is down or the host has failed and cannot be recovered, you can delete its drives by using a different procedure.

See [“Deleting all devices from a media server”](#) on page 392.

To delete a drive

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Select the required drive.
- 4 Select **Delete**.

Note: It may take a few minutes for the web UI to reflect that the drive is deleted. You are prompted to restart the Media Manager device daemon.

Resetting a drive

Resetting a drive changes the state of the drive.

Usually you reset a drive when its state is unknown, which occurs if an application other than NetBackup uses the drive. When you reset the drive, it returns to a known state before use with NetBackup. If a SCSI reservation exists on the drive, a reset operation from the host that owns the reservation can help the SCSI reservation.

If the drive is in use by NetBackup, the reset action fails. If the drive is not in use by NetBackup, NetBackup tries to unload the drive and set its run-time attributes to default values.

Note that a drive reset does not perform any SCSI bus or SCSI device resets.

To reset a drive

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Locate the drive that you want to reset. Then select **Actions > Reset drive**.
- 4 If the drive is in use by NetBackup and cannot be reset, restart the NetBackup Job Manager (`nbjrm`) to free up the drive.
- 5 Determine which job controls the drive (that is, which job writes to or reads from the drive).

On the left, select **Activity monitor**. Then on the **Jobs** tab, cancel the job.

- 6 In the **Activity monitor**, restart the NetBackup Job Manager, which cancels all NetBackup jobs in progress.

Resetting the mount time of a drive

You can reset the mount time of the drive. Reset the mount time to zero after you perform a manual cleaning.

To reset the mount time

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Select a drive.
- 4 Click **Actions > Drive cleaning > Reset mount time**. The mount time for the selected drive is set to zero.
- 5 If you use the Shared drive (Shared Storage Option), do the following:
In the list of hosts that share the drive, choose only one host on which the function applies.
- 6 Click **Reset mount time**.

Setting the drive cleaning frequency

When you add a drive to NetBackup, you configure the automatic, frequency-based cleaning interval. From the **Device monitor** you can change the cleaning frequency that was configured when you added the drive.

To set the cleaning frequency

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Select a drive.
- 4 Click **Actions > Drive cleaning > Set cleaning frequency**.
- 5 Enter the number of mount hours between drive cleaning.

The **Set cleaning frequency** option is not available for the drives that do not support frequency-based cleaning. This function is not available for shared drives.

The drive cleaning interval appears in the **Drive properties**.
- 6 Click **Save**.

Viewing drive details

You can obtain detailed information about drives (or shared drives), such as drive cleaning, drive properties, drive status, host, and robotic library information.

To view the drive details

- 1 Open the web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Many drive details are displayed on this tab. For additional details, select the link for a drive name.

For shared drives, you can see the drive **Control** mode and **Drive index** for each host that shares a drive. Click on the **Shared drive hosts** tab to view a list of hosts that share a drive.

Performing device diagnostics

Diagnostic functions let you run and manage drive and robot diagnostic tests. Diagnostics are executed in an ordered sequence to verify the functionality of hardware devices. These tests can help you to troubleshoot drive or robot problems.

Running a robot diagnostic test

Use this procedure to run diagnostic tests on TLD robotic libraries.

Ensure that the library to be tested is properly configured for use with NetBackup. The existing NetBackup robotic control daemons or processes are used for the test.

Note: NetBackup does not support diagnostic tests for API-attached robotic tape libraries and other types of SCSI-attached libraries.

To run a robot diagnostic test

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Robot Diagnostics**.
- 3 In the **Robot Diagnostics** dialog box, select the media server that is the **Device Host** for the robot that you want to test.
- 4 In the **Robot Name** field, select the robot that you want to diagnose.

- 5 Click **Start** to start the diagnostic tests.

The **Results** window shows results of each step in the test.

Operator intervention is required if the **State** column of the **Results** window contains **Waiting**. For example, a test step may prompt you to load a new tape into a drive before the test can continue.

- 6 If operator intervention is required, select the test step in the **Results** window and click **Details** to determine what you must do. Complete the requested operation task and then click **Continue** in the **Test Details** dialog box to resume the test

To stop a test and change the device

- 1 Click **Stop**.

The test ends after it performs any necessary clean-up work and updates the test records to reflect that the test run has been stopped.

- 2 In the **Device Host** and the **Robot Name** boxes, select the host and the robot that you want to test.
- 3 Click **Start** to restart the diagnostic test.

Running a tape drive diagnostic test

NetBackup diagnostic functions let you run and manage diagnostic tests. Diagnostics are performed in an ordered sequence to verify the functionality of hardware devices. These tests can help you to troubleshoot drive problems.

To run a tape drive diagnostic test

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Drive Diagnostics**.
- 3 In the **Drive Diagnostics** dialog box, select the media server that contains the drive that you want to test in the **Device Host** box.
- 4 In the **Drive Name** box, select the drive.

- 5 Click **Start** to start the diagnostic tests.

For robotic drives, the test media is loaded automatically.

For a standalone drive, insert the prelabeled test tape that is shown in the **Step Information** column of the **Results** window.

The **Results** window shows results of each step in the test.

- 6 If operator intervention is required, the State column of the Results window displays Waiting. For example, a test step may require that you to load a new tape into a drive before the test can continue.

Complete the intervention and then click **Continue**.

Select the test step in the **Results** window and click **Details** to determine what you must do. Complete the requested operation task and then click **Continue** in the **Test Details** dialog box to resume the test

To stop a test and change the device

- 1 Click **Stop**.

The test ends after it performs any necessary clean-up work and updates the test records to reflect that the test run has been stopped.

- 2 In the **Device Host** and the **Drive** boxes, select the host and the drive that you want to test.

- 3 Click **Start** to restart the diagnostic test.

Managing a diagnostic test step that requires operator intervention

Operator intervention is required if the **Status** column of the **Results** display contains **Waiting**. For example, a test step may prompt for a new tape to be loaded into a drive before the test continues.

To manage a diagnostic step

- 1 Complete the requested operations task.
- 2 Click **Continue** to resume the test.

If you clicked **Details** for a test step that requires operator intervention, you can click **Continue** from the **Test Details** dialog box.

Obtaining detailed information for a diagnostic test step

You can get information for a test step at any time during the test.

To obtain detailed information for a diagnostic test step

- 1 Select a test step in the **Results** display.
- 2 Click **Details**. A dialog box appears that displays information for the step.

The information includes a brief explanation of the checks that are performed by a specific step and the instructions that are associated with any step that requires manual intervention. For example, a step may prompt for a new tape to be loaded into a tape drive before the diagnostic session continues.
- 3 Click **Close** to return to the **Device Diagnostics** dialog box.

Verifying the device configuration

Verify the device configuration by running the **Configure drives and robots** wizard. However, some details of a device configuration cannot be validated without attempting tape mounts. Use the NetBackup `robtest` utility to mount tapes and validate the configuration.

To verify robots and drives by using the wizard

- 1 On the left, select **Storage > Tape storage**.
- 2 On the top right, select **Actions > Configure storage devices** and follow the wizard instructions.

About automatic path correction

NetBackup automatic path correction recognizes if you change a device because the serial number of the new device is different than the serial number of the old device. NetBackup updates the device configuration automatically.

NetBackup recognizes device changes as follows:

- When the NetBackup Device Manager (`ltid`) performs automatic path correction. See [“About the NetBackup Device Manager”](#) on page 460.
- When the Windows Plug-n-Play feature performs serial number checks.

By default, Windows and Linux systems are configured for automatic path correction. On other operating systems, you must enable it.

See [“Enabling automatic path correction”](#) on page 457.

In some circumstances, NetBackup may be unable to determine the correct serial number in a small number of tape drives and robotic libraries. For example, NetBackup may configure serialized devices as unserialized or configure a device

with the wrong serial number. If so, a device may be unusable (such as the tape drive may be downed).

To resolve such a problem, do one of the following actions:

- Configure the new device with the **Configure drives and robots** wizard.
The server operating system must recognize the device before you can configure it in NetBackup. Device configuration can require remapping, rediscovery, and possibly a restart of the operating system.
For more information, see the [NetBackup Device Configuration Guide](#).
- Disable the automated device discovery by using the `vm.conf` file `AUTO_PATH_CORRECTION` option.

Enabling automatic path correction

You can enable automatic device path correction in NetBackup. To do so, use the following procedure.

See [“About automatic path correction”](#) on page 456.

To configure automatic path correction

- 1 Use a text editor to open the following file:

On Windows:

```
install_path\Cohesity\Volmgr\vm.conf
```

On UNIX:

```
/usr/opensv/volmgr/vm.conf
```

- 2 Add the following `AUTO_PATH_CORRECTION` entry to the file:

```
AUTO_PATH_CORRECTION = YES
```

If it already exists but is set to **NO**, change the value to **YES**.

- 3 Save the file and exit the text editor.

Replacing a device

Two processes exist for replacing a device, as follows:

Replace a device on a single host See [Table 8-8](#) on page 458.

Replace a shared device See [Table 8-9](#) on page 458.

Table 8-8 To replace a device on a single host

Task	Instructions
If the device is a drive, change the drive state to DOWN.	See “Changing a drive operating mode” on page 446.
Replace the device. Specify the same SCSI ID for the new device as the old device.	See the vendor's documentation.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 446.
If either of the following are true, configure the new device by using the Configure drives and robots wizard: ■ You replaced a drive with a different drive type. ■ You replaced a serialized drive with an unserialized drive.	See the NetBackup Web UI Administrator's Guide .

Table 8-9 To replace a shared device

Task	Instructions
If the device is a drive, change the drive state to DOWN.	See “Changing a drive operating mode” on page 446.
Replace the device. Specify the same SCSI ID for the new device as the old device.	See the vendor's documentation.
Produce a list of new and missing hardware.	The following command scans for new hardware and produces a report that shows the new and the replaced hardware: On Windows: <pre>install_path\Cohesity\Volmgr\bin\tpautoconf -report_disc</pre> On UNIX: <pre>/usr/opensv/volmgr/bin/tpautoconf -report_disc</pre>
Ensure that all servers that share the new device are up and that all NetBackup services are active.	See the NetBackup Web UI Administrator's Guide for information on how to monitor NetBackup daemons.

Table 8-9 To replace a shared device (*continued*)

Task	Instructions
Read the serial number from the new device and update the NetBackup database.	If the device is a robot, run the following command: On Windows: <pre>install_path\Cohesity\Volmgr\bin\tpautoconf -replace_robot robot_number -path robot_path</pre> On UNIX: <pre>/usr/opensv/volmgr/bin/tpautoconf -replace_robot robot_number -path robot_path</pre> If the device is a drive, run the following commands: On Windows: <pre>install_path\Cohesity\Volmgr\bin\tpautoconf -replace_drive drive_name -path path_name</pre> On UNIX: <pre>/usr/opensv/volmgr/bin/tpautoconf -replace_drive drive_name -path path_name</pre>
If the new device is an unserialized drive, run the Configure drives and robots on all servers that share the drive. If the new device is a robot, run the Configure drives and robots on the server that is the robot control host.	See the NetBackup Web UI Administrator's Guide .
If the device is a drive, change the drive state to UP.	See "Changing a drive operating mode" on page 446.

Updating device firmware

By default, NetBackup recognizes if you update the firmware of a device.

The following table describes an overview of how to update device firmware.

Table 8-10 How to update device firmware

Task	Instructions
If the device is a drive, change the drive state to DOWN.	See "Changing a drive operating mode" on page 446.

Table 8-10 How to update device firmware (*continued*)

Task	Instructions
Update the firmware.	See the vendor's documentation.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 446.

About the NetBackup Device Manager

The NetBackup Device Manager (`ltid`) manages robot and tape processes for NetBackup. The Device Manager processes requests to mount and unmount tapes in robotically controlled devices through the robotic control processes. NetBackup starts `ltid` on the hosts that have storage devices configured. The Device Manager starts the Volume Manager (`vmd`), the automatic volume recognition process (`avrd`), and any robotic processes as needed.

In the NetBackup web UI, the NetBackup Device Manager is exposed as follows:

In the Activity monitor	As <code>ltid</code> .
In Storage > Tape storage	On the Actions menu, the menu item Stop/Restart media manager device daemon .

Note: If you stop and restart the Device Manager, any backups, archives, or restores that are in progress may fail.

About external access to NetBackup controlled devices on UNIX

On UNIX hosts, the NetBackup Device Manager restricts access to drives that are in an `UP` state by changing the permissions of the device files for those drives. The Device Manager changes the permissions to 0600 when it starts and back to their original settings when it is terminated. The permissions also are returned to their original settings when a drive's state is changed to `DOWN`.

See [“About the NetBackup Device Manager”](#) on page 460.

Do not modify the permissions of these device files when the Device Manager is active. The automatic volume recognition process (`avrd`) periodically tries to rewind and read data from media in the drives that are `UP` and are not currently assigned in NetBackup.

To ensure reliable operation, do not use UNIX tape and drive commands on the drives that are `UP` and controlled by the Device Manager. Users can use the NetBackup `tpreq` and `tpunmount` commands and the `drive_mount_notify` and `drive_unmount_notify` scripts on those drives.

For more information, see “NetBackup notify scripts” in the [NetBackup Administrator's Guide, Volume II](#).

Stopping and restarting the device manager

Use the following procedure to stop and restart the NetBackup Device Manager.

See [“About the NetBackup Device Manager”](#) on page 460.

To start or stop the Device Manager

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Stop/Restart Media Manager Device Daemon**.
- 3 In the **Stop/Restart Media Manager Device Daemon** dialog box, do the following:
 - 1 In the **Device host** list, select the media server that you want to restart.
 - 2 Select the **Action: Start, Stop, or Stop/Restart**.
The actions that are available depend on the state of the device manager or daemon.
 - 3 Select the wanted **Options: Eject media from standalone drive(s)** or **Enable verbose logging**.
 - 4 Click **Apply** or **OK**, depending on the following results:
 - **Apply** does not close the dialog box so you can select device hosts and actions for more than another device host.
 - **OK** closes the dialog box.

Configuring tape media

This chapter includes the following topics:

- [About NetBackup tape volumes](#)
- [About NetBackup volume pools](#)
- [About NetBackup volume groups](#)
- [NetBackup media types](#)
- [About WORM media](#)
- [About adding volumes](#)
- [About configuring media name and attribute rules](#)
- [Adding volumes](#)
- [Volume properties](#)
- [About media settings rules](#)
- [Configuring media settings](#)
- [About barcodes](#)
- [Configuring barcode rules](#)
- [About media ID generation rules](#)
- [Configuring media ID generation rules](#)
- [About media type mapping rules](#)
- [Configuring media type mappings](#)
- [Managing volumes](#)

- [Managing volume pools](#)
- [Managing volume groups](#)
- [About media sharing](#)
- [Configure unrestricted media sharing](#)
- [Configure media sharing with a server group](#)

About NetBackup tape volumes

A tape volume is a data storage tape or a cleaning tape. NetBackup assigns attributes to each volume and uses them to track and manage the volumes. Attributes include the media ID, robot host, robot type, robot number, and slot location.

NetBackup uses two volume types, as follows:

Robotic volumes	Volumes that are located in a robot. The robotic library moves the volumes into and out from the robotic drives as necessary.
Standalone volumes	Volumes that are allocated for the drives that are not in a robot. Operator intervention is required to load volumes into and eject volumes from standalone drives.

NetBackup uses volume pools to organized volumes by usage.

See [“About NetBackup volume pools”](#) on page 463.

Volume information is stored in the NetBackup database.

About NetBackup volume pools

A volume pool identifies a set of volumes by usage. Volume pools protect volumes from access by unauthorized users, groups, or applications. When you add media to NetBackup, you assign them to a volume pool (or assign them as standalone volumes, without a pool assignment).

By default, NetBackup creates the following volume pools:

NetBackup	The default pool to which all backup images are written (unless you specify otherwise).
DataStore	For DataStore use.

CatalogBackup	For NetBackup catalog backups. Catalog backup volumes are not a special type in NetBackup. They are the data storage volumes that you assign to the CatalogBackup volume pool. To add NetBackup catalog backups, use any of the add volume methods. Ensure that you assign them to the volume pool you use for catalog backups. After adding volumes, use the NetBackup Catalog Backup wizard to configure a catalog backup policy.
None	For the volumes that are not assigned to a pool.

You can add other volume pools. For example, you can add a volume pool for each storage application you use. Then, as you add volumes to use with an application, you assign them to that application's volume pool. You can also move volumes between pools.

You also can configure a scratch pool from which NetBackup can transfer volumes when a volume pool has no volumes available.

The volume pool concept is relevant only for tape storage units and does not apply to disk storage units.

You can use any of the approved characters for volume pool names.

NetBackup uses several special prefixes for volume pool names.

Examples of volume pool usage are available in the [NetBackup Administrator's Guide, Volume II](#).

About reserved volume pool name prefixes

NetBackup reserves the following prefixes (case sensitive) for the names of the volume pools that contain media for specific purposes:

ENCR	For volumes on which NetBackup encrypts the data. The volumes in a volume pool that uses this name prefix must be targeted to encrypting tape drives. See the NetBackup Security and Encryption Guide .
WENCR	For WORM volumes on which NetBackup encrypts the data. The volumes in a volume pool that uses this name prefix must be targeted to encrypting tape drives. See " About NetBackup encryption options " on page 691. See the NetBackup Security and Encryption Guide . See the NetBackup Commands Reference Guide

WORM For WORM volumes. NetBackup does not encrypt the data.
See [“About using volume pools to manage WORM media”](#) on page 469.

NetBackup examines the volume pool names to determine if they are special purpose volume pools. If a volume pool name begins with one of the special prefixes, NetBackup processes the media in that pool according to the requirements for that pool. If not, NetBackup does not use special processing for that volume pool's media.

When you create a volume pool for any of these purposes, you must use uppercase characters. For readability, it may be beneficial to use an underscore character after the prefix, for example `WORM_` or `ENCR_`.

About scratch volume pools

The scratch pool is an optional pool that contains the media that NetBackup can allocate to other pools as needed. If you configure a scratch pool, NetBackup moves volumes from that scratch pool to other pools that do not have volumes available.

Only one scratch pool is allowed. You cannot add a scratch pool if one exists.

You cannot change the **NetBackup** or **DataStore** pools to be scratch volume pools.

If you create a scratch pool, be aware of the following conditions:

- If the scratch pool contains assigned volumes, these volumes remain in the scratch pool.
NetBackup does not move assigned volumes to other pools as it does with unassigned volumes.
- NetBackup does not assign volumes while they are in a scratch pool.
For example if a NetBackup policy or schedule specifies the scratch pool, all requests for those volumes are denied.
- NetBackup returns expired media to the scratch volume pool automatically (media that is returned must have been originally in the same scratch pool).
- To use NetBackup to manage the allocation of volumes to volume pools, do the following:
 - Create volume pools as required, but do not add any volumes to the pools.
 - Define a scratch pool and add all of the volumes to it. NetBackup moves volumes to the other pools as volumes are needed.

See [“About NetBackup volume pools”](#) on page 463.

See [“About media settings rules”](#) on page 478.

About NetBackup volume groups

A volume group identifies a set of volumes that reside at the same physical location. The location can be either the robot in which the volumes reside, standalone storage, or off-site storage if you use the NetBackup Vault option.

When you add media to NetBackup, NetBackup assigns all volumes in a robot to that robot's volume group. Alternatively, you can assign the media to a different group.

Volume groups are convenient for tracking the location of volumes, such as the case when a volume is moved off site. Volume groups let you perform operations on a set of volumes by specifying the group name rather than each individual media ID of each volume. Operations include moves between a robotic library and a standalone location or deletions from NetBackup.

If you move a volume physically, you also must move it logically. A logical move means to change the volume attributes to show the new location.

The following are the rules for assigning volume groups:

- All volumes in a group must be the same media type.
However, a media type and its corresponding cleaning media type are allowed in the same volume group (such as DLT and DLT_CLN).
- All volumes in a robotic library must belong to a volume group.
You cannot add volumes to a robotic library without specifying a group or having Media Manager generate a name for the group.
- The only way to clear a volume group name is to move the volume to standalone and not specify a volume group.
- More than one volume group can share the same location.
For example, a robotic library can contain volumes from more than one volume group and you can have more than one standalone volume group.
- All volumes in a group must be in the same robotic library or be standalone.
That is, you cannot add a group (or part of a group) to a robotic library if it already exists in another robotic library.

Examples of volume group usage are available.

Examples of volume group usage are available in the [NetBackup Administrator's Guide, Volume II](#).

NetBackup media types

NetBackup uses media types to differentiate the media that have different physical characteristics. Each media type may represent a specific physical media type.

The NetBackup media types are also known as Media Manager media types. The following table describes the NetBackup media types.

Table 9-1 NetBackup media types

Media type	Description
DLT	DLT cartridge tape
DLT_CLN	DLT cleaning tape
DLT2	DLT cartridge tape 2
DLT2_CLN	DLT cleaning tape 2
DLT3	DLT cartridge tape 3
DLT3_CLN	DLT cleaning tape 3
HCART	1/2 inch cartridge tape
HCART2	1/2 inch cartridge tape 2
HCART3	1/2 inch cartridge tape 3
HC_CLN	1/2 inch cleaning tape
HC2_CLN	1/2 inch cleaning tape 2
HC3_CLN	1/2 inch cleaning tape 3

NetBackup writes media in a format that allows the position to be verified before NetBackup appends new backup images to the media.

Note: The user interface for NetBackup may show configuration options for the media types that are not supported in that release. Those types may be supported in an earlier release, and a NetBackup primary server can manage the hosts that run earlier NetBackup versions. Therefore, the configuration information for such types must appear in the user interface. The NetBackup documentation also may describe the configuration information for such types. To determine which versions of NetBackup support which media types, see the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List*:

https://www.veritas.com/support/en_US/article.100040093

Alternate NetBackup media types

Alternate media types let you define more than one type of tape in the same library. You can use the alternate types to differentiate between different physical cartridges.

The following are examples of alternate media types:

- DLT, DLT2, DLT3
- HCART, HCART2, HCART3

For example, if a robot has DLT4000 and DLT7000 drives, you can specify the following media types:

- DLT media type for the DLT4000 tapes
- DLT2 media type for the DLT7000 tapes

NetBackup then does not load a tape that was written in a DLT4000 drive into a DLT7000 drive and vice versa.

You must use the appropriate default media type when you configure the drives. (When you configure drives in NetBackup, you specify the default media type to use in each drive type.)

In a robot, all of the volumes (of a specific vendor media type) must be the same NetBackup media type. For example, for an ACS robot that contains 3490E media, you can assign either NetBackup HCART, HCART2, or HCART3 media type to that media. You cannot assign HCART to some of the media and HCART2 (or HCART3) to other of the media.

For more information, see "Media formats" in the [NetBackup Administrator's Guide, Volume II](#).

About WORM media

You can use WORM (Write-Once-Read-Many) media to protect key data from unwanted modification or to meet compliance regulations.

NetBackup uses the QIC/WORM tape format for WORM media. This format lets NetBackup append images to WORM tape.

For more information about "Media formats", see the [NetBackup Administrator's Guide, Volume II](#).

Tape error recovery is disabled for WORM media. NetBackup has job resume logic, which tries to resume a job that has been interrupted (such as an interruption on the Fibre Channel). However, NetBackup fails a job that uses WORM media and then retries the failed job. It is recommended that you use checkpoint and restart for backups.

The `bplabel` command labels only LTO-3 WORM tapes. All other WORM media cannot be labeled because the label cannot be overwritten when the media is used.

The following are the limitations for WORM tape:

- Third-party copy backups are not supported with WORM media.
- NetBackup does not support resume logic with WORM tape. NetBackup fails a job that uses WORM media and then retries the failed job. Alternatively, if checkpoint and restart are used, NetBackup restarts the job from the last checkpoint. It is recommended that you use checkpoint and restart for backups.

NetBackup provides two methods to manage WORM media, as follows:

- Assign a reserved prefix to WORM volume pool names.
A WORM volume name cannot contain a period (.).
See [“About using volume pools to manage WORM media”](#) on page 469.
- Assign a specific drive type to all WORM drives and a specific media type to all WORM media.
See [“About using unique drive and media types to manage WORM media”](#) on page 470.

About using volume pools to manage WORM media

You can dedicate volume pools for WORM media. This method lets a WORM-capable tape drive back up and restore standard and WORM media. NetBackup uses two reserved volume pool prefixes to indicate that the volumes in a pool are for WORM drives, as follows:

- WORM (uppercase letters) denotes WORM media.
- WENCR (uppercase letters) denotes WORM media on which NetBackup should encrypt the data.

See [“About reserved volume pool name prefixes”](#) on page 464.

For more information about encrypting data on your media, see the [NetBackup Security and Encryption Guide](#).

When you create a volume pool for WORM media, specify one of the reserved prefixes as the first characters of the pool name. NetBackup examines the volume pool names to determine if they begin with a reserved prefix. For readability, it may be beneficial to use an underscore character after the prefix, for example **WORM_**.

See [“Add a volume pool”](#) on page 521.

See [“Edit or delete a volume pool”](#) on page 521.

Note the following cases:

- If the drive contains WORM media and the media is in a WORM volume pool, NetBackup writes the media as WORM.
- If the drive contains WORM media and the media is not in a WORM volume pool, NetBackup freezes the media.
- If the drive contains standard media and the media is in a WORM volume pool, NetBackup freezes the media.
- If the drive contains the Quantum media that has never been used or all of its NetBackup images have expired, NetBackup uses the media.

See [“About using a WORM scratch pool”](#) on page 470.

See [“About WORM media”](#) on page 468.

See [“About using unique drive and media types to manage WORM media”](#) on page 470.

About using a WORM scratch pool

For all supported WORM-capable drives (except the Quantum drive), the scratch pool must only contain one type of media. It is recommended that you add the most commonly used media to the scratch pool. For example, if most NetBackup jobs use standard media, put standard media in the scratch pool.

If the scratch pool contains standard media, ensure that the WORM volume pool does not run out of media to complete backup jobs.

If the WORM volume pool runs out of media, NetBackup performs the following actions:

- Moves the standard media from the scratch pool into the WORM pool.
- Loads the standard media into a WORM-capable drive.
- Freezes the media.

NetBackup repeats this process until all of the standard media in the scratch pool is frozen.

The opposite also is true. If a standard volume pool runs out of media and the scratch pool contains WORM media, standard backups can fail because appropriate media are unavailable.

About using unique drive and media types to manage WORM media

You can assign a different drive and media type to all WORM drives and media. For example, configure standard drives and media as HCART and WORM-capable drives and media as HCART2.

This method lets you add both types of media in the scratch pool because NetBackup selects the correct media type for the drive type.

However, because each drive is limited to backups and restores with a specific type of media, optimal drive usage may not be achieved. For example, the WORM-capable drives cannot be used for backups with standard media even if no WORM backups are in progress.

Because Quantum drives use only a single media type, this method for managing the WORM media is unnecessary.

See [“About WORM media and the Quantum drive”](#) on page 471.

If you use unique drive and media types to manage WORM media, disable the WORM volume pool name verification.

See [“Disabling WORM volume pool name verification”](#) on page 471.

Disabling WORM volume pool name verification

If you use unique drive and media types to manage WORM media, disable NetBackup volume pool name verification. WORM volume pool name verification is used only for the WORM volume pool method of managing WORM media.

See [“About using unique drive and media types to manage WORM media”](#) on page 470.

See [“About using volume pools to manage WORM media”](#) on page 469.

To disable the volume pool name verification

- ◆ Create the following touch file on the media server of the WORM drive:

On Windows:

```
install_path\NetBackup\db\config\DISABLE_WORM_POOLCHECK
```

On UNIX:

```
/usr/opensv/netbackup/db/config/DISABLE_WORM_POOLCHECK
```

About WORM media and the Quantum drive

When you use the Quantum drive, only one kind of media can be used as either standard media or WORM media.

If a WORM volume pool runs out of media, media are moved from the scratch volume pool into the WORM pool. NetBackup determines whether the media are configured as standard or WORM media. For a standard media volume, NetBackup reads the tape label and verifies that the media is unused or that all images are expired. NetBackup also verifies that the media is not currently assigned to a server.

After verification, NetBackup configures the media as WORM media and continues with the NetBackup job.

Supported WORM drives

NetBackup requires a SCSI pass-through driver to use WORM tape drives. NetBackup queries the drive to verify that drive is WORM-capable and that the media in the drive is WORM media. SCSI pass-through paths are provided on the server platforms NetBackup supports. SCSI pass-through paths may require special operating system configuration changes.

For information about the drives that NetBackup supports for WORM media, see the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List* at the following URL:

https://www.veritas.com/support/en_US/article.100040093

All of the vendors except Quantum require the use of special WORM media.

Quantum lets NetBackup convert standard tape media to WORM media. To use Quantum drives for WORM media on Solaris systems, modify the `st.conf` file.

For more Information about how to configure nonstandard tape drives and how to edit the `st.conf` file, see the [NetBackup Device Configuration Guide](#).

About adding volumes

Adding volumes is a logical operation that assigns NetBackup attributes to physical media. The media can reside in storage devices already, or you can add them to the storage devices when you add them to NetBackup. How you add volumes depends on the type of volume: robotic or standalone.

See [“About adding robotic volumes”](#) on page 472.

See [“About adding standalone volumes”](#) on page 473.

NetBackup uses the rules to assign names and attributes to volumes.

About adding robotic volumes

The robotic volumes are the volumes that are located in a robotic tape library. The following table describes the methods for adding robotic volumes.

Table 9-2 Methods for adding robotic volumes

Method	Description
Manually add a volume	See “Adding volumes” on page 474.

Table 9-2 Methods for adding robotic volumes (*continued*)

Method	Description
Robot inventory	See “About robot inventory” on page 528. See “About updating the NetBackup volume configuration” on page 537.
NetBackup commands	See the NetBackup Commands Reference Guide .

About adding standalone volumes

Standalone volumes are the volumes that reside in the drives that are not in a robot or are allocated for standalone drives.

Because NetBackup does not label volumes until it uses them, you can add volumes even though they do not reside in a drive. The additional volumes are available for use if the volume in a drive becomes full or unusable. For example, if a volume in a standalone drive is full or unusable because of errors, NetBackup ejects (logically) the volume. If you add other standalone volumes, NetBackup requests that volume; NetBackup does not generate an `out of media` error.

The `DISABLE_STANDALONE_DRIVE_EXTENSIONS` option of the `nbemmcmd` command can turn off the automatic use of standalone volumes.

Table 9-3 Methods for adding standalone volumes

Method	Description
Manually add a volume	See “Adding volumes” on page 474.
NetBackup commands	See the NetBackup Commands Reference Guide .

About configuring media name and attribute rules

NetBackup uses the default settings and rules to name and assign attributes to new removable media. NetBackup uses these rules when you do the following:

- Add new media.
To add new media, run the **Configure drives and robots** wizard. On the left select **Storage > Tape storage**. Then select **Actions > Configure storage devices**.
- Inventory a robot.

To inventory a robot, on the left select **Storage > Tape storage**. Then select the **Robots** tab. Select the robot, then select **Inventory robot**. If NetBackup discovers new media in the robot, it adds that media to NetBackup.

For most configurations, the default settings work well. However, you can change the default settings and rules that NetBackup uses. Change the settings only if you have special hardware or usage requirements. You can change the settings from the **Configure drives and robots** or in the **Inventory robot** settings.

The following table shows the rules that you can configure:

Table 9-4 Attributes for media

What	Where
Media settings	See “About media settings rules” on page 478. See “Configuring media settings” on page 478.
Barcode rules	See “About barcodes” on page 486. See “Configuring barcode rules” on page 490.
Media ID generation rules	See “About media ID generation rules” on page 492. See “Configuring media ID generation rules” on page 493.
Map media for API robots	See “About media type mapping rules” on page 496. See “Configuring media type mappings” on page 496.

Adding volumes

Use this procedure to add a new volume.

Be careful when you specify properties. You cannot change some properties later, such as the media ID or type. If you specify them incorrectly, you must delete the volume and add it again.

To add a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select **Add volume**.

5 Specify the properties for the volumes.

The properties that display can vary depending on the type of volume.

See “[Volume properties](#)” on page 475.

6 Select **Save**.

If the robot has a barcode reader, NetBackup performs the following actions:

- Adds the volume to the EMM database using the specified media ID.
- Reads the barcode of each new volume.
- Adds the barcodes as attributes in the EMM database.

Volume properties

Volume properties describes the properties for volumes in NetBackup. The properties depend on whether you add, change, or move volumes.

The properties are arranged alphabetically.

Table 9-5 Volume properties

Property	Description	Operation
Device host	The name of the NetBackup media server to which the robot is attached.	Add, move
Expiration date	The following does not apply to cleaning tapes. The date after which the volume is too old to be reliable. When the expiration date has passed, NetBackup reads data on the volume but does not mount and write to the volume. You should exchange it for a new volume. When you add a new volume, NetBackup does not set an expiration date. The expiration date is not the same as the retention period for the backup data on the volume. You specify data retention periods in the backup policies.	Change
First media ID	This property appears only if the number of volumes is more than one. The ID of the first volume in the range of volumes. Media IDs need to be exactly 6 characters. Valid only when you add a range of volumes. Use the same pattern that you chose in the Media ID naming style box. NetBackup uses the pattern to name the remaining volumes by incrementing the digits. NetBackup allows specific characters in names.	Add

Table 9-5 Volume properties (*continued*)

Property	Description	Operation
First slot number	The number of the first slot in the robot in which the range of volumes resides. If you add or move more than one media, NetBackup assigns the remainder of the slot numbers sequentially. Note: You cannot enter slot information for volumes in an API robot. The robot vendor tracks the slot locations for API robot types.	Add, move
Maximum cleanings	The maximum number of times NetBackup should mount the volume or use the cleaning tape. To determine the maximum mount limit to use, consult the vendor documentation for information on the expected life of the volume.	Add
Maximum mounts	The following topic does not apply to cleaning tapes. The Maximum mounts property specifies the number of times that the selected volumes can be mounted. When the limit is reached, NetBackup reads data on the volume but does not mount and write to the volume. A value of zero (the default) is the same as Unlimited. To help determine the maximum mount limit, consult the vendor documentation for information on the expected life of the volume.	Add, change
Media description	A description of the media, up to 25 character maximum. NetBackup allows specific characters in names.	Add, change
Media ID	This property appears only if the number of volumes is one. The ID for the new volume. Media IDs must be exactly 6 characters. Media IDs for an API robot must match the barcode on the media (for API robots, NetBackup supports barcodes of 6 characters). Therefore, obtain a list of the barcodes before you add the volumes. Obtain this information through a robotic inventory or from the robot vendor's software. NetBackup allows specific characters in names.	Add, change
Media ID naming style	The style to use to name the range of volumes. Media IDs must be exactly 6 characters in length. Using the pattern, NetBackup names the remaining volumes by incrementing the digits. NetBackup media IDs for an API robot must match the barcode on the media. For API robots, NetBackup supports barcodes from 1 to 6 characters. Therefore, obtain a list of the barcodes before you add the volumes. Obtain this information through a robotic inventory or from the robot vendor's software. NetBackup allows specific characters in names.	Add

Table 9-5 Volume properties (*continued*)

Property	Description	Operation
Media type	The media type for the volume to add. Select the type from the drop-down list.	Add
Number of volumes	The number of volumes to add. For a robotic library, enough slots must exist for the volumes.	Add
Robot	The robotic library to add or move the volumes to. To add volumes for a different robot, select a robot from the drop-down list. The list shows robots on the selected host that can contain volumes of the selected media type.	Add, move
Volume group	If you specified a robot, select from a volume group already configured for that robot. Alternatively, enter the name for a volume group; if it does not exist, NetBackup creates it and adds the volume to it. If you do not specify a volume group (you leave the volume group blank), the following occurs: ■ Standalone volumes are not assigned to a volume group. ■ NetBackup generates a name for robotic volumes by using the robot number and type. For example, if the robot is a TLD and has a robot number of 50, the group name is 000_00050_TLD. See “About NetBackup volume groups” on page 466. See “About rules for moving volumes between groups” on page 504.	Add, move
Volume is in a robotic library	When you add a volume: ■ If the volume is in a robot, select Volume is in a robotic library. ■ If the volume is a standalone volume, do not select Volume is in a robotic library. When you move a volume: ■ To inject a volume into a robotic library, select Volume is in a robotic library. Then, select a robot and the slot number (First slot number) for the volume. ■ To eject a volume from a robot, clear Volume is in a robotic library.	Add, move

Table 9-5 Volume properties (*continued*)

Property	Description	Operation
Volume pool	The pool to which the volume or volumes should be assigned. Select a volume pool you created or one of the following standard NetBackup pools: ■ None. ■ NetBackup is the default pool name for NetBackup. ■ DataStore is the default pool name for DataStore. ■ CatalogBackup is the default pool name used for NetBackup catalog backups of policy type NBU-Catalog. When the images on a volume expire, NetBackup returns it to the scratch volume pool if it was allocated from the scratch pool. See “About NetBackup volume pools” on page 463.	Add, change
Volumes to move	The Volumes to move section of the dialog box shows the media IDs of the volumes that you selected to move.	Move

About media settings rules

The NetBackup media settings rules depend on the following:

- For existing media, the volume group to which the volumes belong.
- For new media, the media ID prefix, the media type, and the pool to which the volume should be assigned.

You can change the default rules.

See [“Configuring media settings”](#) on page 478.

See [“Media settings options”](#) on page 479.

Configuring media settings

This procedure describes how to configure the attributes for existing and new media.

See [“About media settings rules”](#) on page 478.

To configure media settings

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**.
- 3 On the top right, click **Actions > Inventory robots**.

- 4 Select the **Device host**.
- 5 Select the robot.
- 6 Click either **Preview volume configuration changes** or **Update volume configuration**.
- 7 Click on **Advanced options**. Then click **Media settings**.
- 8 Configure the settings as follows:
 - a. In the **Media which have been removed from the robot should be assigned to the volume group** list, select a volume group for the media that are removed from the robot.
 - b. In the **Media which have been moved into or within the robot should be assigned to the volume group** list, select a volume group for the media that are in or are added to the robot.
 - c. If the robotic library supports barcodes and the volume has readable barcodes, NetBackup creates media IDs automatically from the barcodes. You do not need to configure a prefix.

However, if the media in the robotic library has unreadable barcodes or if the robot does not support barcodes, NetBackup assigns a default media ID prefix.

To use a media ID prefix other than the **Default**, click **Use the following Media ID prefix** field. Then, specify or choose a media ID prefix.
 - d. To use your barcode rules to assign attributes to new volumes, select **Use barcode rules**.
 - e. To override your barcode rules for the new media in the robotic library, select a **Media type** from the list.
 - f. To override the default volume pool for the new media in the robotic library, select a **Volume pool** from the list.
- 9 Click **Save**.

Media settings options

The following are the settings for the new media in a robot that you add to your NetBackup volume configuration.

Media which have been removed from the robot... (existing media setting)

For the media that already exist in your volume configuration, you can specify the volume group if the media are removed from the robot. **Media which have been removed from the robot should be assigned to the volume group.**

The **Media which have been removed from the robot should be assigned to the volume group** drop-down box contains the following selections:

AUTO GENERATE NetBackup automatically generates a new volume group.

DEFAULT If there is an existing group with a compatible residence for the volume, the volume is added to that group. If a suitable volume group does not exist, NetBackup generates a new volume group name.

NO VOLUME GROUP The media are not assigned to a volume group.

Other selections may be available, depending on the setting of the **Media type** field of the **New media** section of the dialog box, as follows:

If the **Media type** field is **DEFAULT** The **Media which have been removed from the robot should be assigned to the volume group** dropdown box includes the volume groups that are valid for the robot's default media type.

If the **Media type** field is other than **DEFAULT** The **Media which have been removed from the robot should be assigned to the volume group** dropdown box includes the volume groups that are valid for the specified media type.

To specify a volume group other than **DEFAULT**, enter a volume group name or select one from the list.

See [“Media type \(new media setting\)”](#) on page 483.

Media which have been moved into or within the robot... (existing media setting)

You can specify the volume group for the existing media that have been moved into or within a robot.

The volume group to assign to the existing media that you have inserted into the robot (or moved to a new location within the robot).

The **Media which have been moved into or within the robot should be assigned to the volume group** drop-down box contains the following selections:

AUTO GENERATE NetBackup automatically generates a new volume group.

DEFAULT If there is an existing group with a compatible residence for the volume, the volume is added to that group. If a suitable volume group does not exist, NetBackup generates a new volume group name.

The following other selections may be available depending on the setting of the **Media type** field of the **New media** section of the dialog box:

If the **Media type** field is **DEFAULT** The **Media which have been moved into or within the robot should be assigned to the volume group** drop-down box includes the volume groups that are valid for the robot's default media type.

If the **Media type** field is other than **DEFAULT** The **Media which have been moved into or within the robot should be assigned to the volume group** drop-down box includes the volume groups that are valid for the specified media type.

To specify a volume group other than **DEFAULT**, enter a volume group name or select one from the list.

If the robotic library contains multiple media types, a **DEFAULT** setting is recommended. If you specify a volume group and volumes of different media types were moved into or within the robot, the new update fails. Volumes of different media types cannot have the same volume group.

See [“Media type \(new media setting\)”](#) on page 483.


Use the following Media ID prefix (new media setting)

if the media has unreadable barcodes or if the robot does not support barcodes, by default NetBackup assigns media ID prefixes as follows:

- If **MEDIA_ID_PREFIX** entries are defined in the `vm.conf` file, NetBackup assigns the last **MEDIA_ID_PREFIX** entry as the media ID prefix.
- If no **MEDIA_ID_PREFIX** entries are defined in the `vm.conf` file, NetBackup uses the letter A as the media ID prefix.

To configure NetBackup to use a media ID prefix other than the default, select **Use the following Media ID prefix** field in the **Advanced Robot Inventory Options** dialog box and then click **Browse**. The **Media ID Prefix** dialog box appears.

Figure 9-1 Media ID Prefix dialog box



The following are the alternative NetBackup media ID assignment behaviors that you can configure in the dialog box:

To not use a media ID prefix	Deselect the Use a media ID prefix for media with unreadable barcodes or if the robot does not support barcodes option.
To use a media ID prefix	Select the Use a media ID prefix for media with unreadable barcodes or if the robot does not support barcodes option.
To use a specific media ID prefix for the current operation only	Select the Specify the media ID prefix for the current session only option then enter the media ID prefix. You can specify a prefix of one to five alphanumeric characters. NetBackup assigns the remaining numeric characters to create a six character media ID. NetBackup uses the prefix only for the current operation.
To configure a media ID prefix to use for the current session and future sessions	Select the Choose from the Media ID prefix list (stored in vm.conf file) option and then select the prefix from the list.
To add a new media ID prefix to the <code>vm.conf</code> file	Select the Choose from the Media ID prefix list (stored in vm.conf file) option and then enter the prefix in the New media ID prefix field. Click Add .

To remove a media ID prefix from the `vm.conf` file Select the **Choose from the Media ID prefix list (stored in `vm.conf` file)** option, select the prefix from the list, and then click **Remove**.

For more information, about the `vm.conf` file, see the [NetBackup Administrator's Guide, Volume II](#).

Use barcode rules (new media setting)

Specifies whether or not to use barcode rules to assign attributes for new media.

To enable barcode rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

See [“About barcodes”](#) on page 486.

See [“Configuring barcode rules”](#) on page 490.

For more information about the `vm.conf` file, see the [NetBackup Administrator's Guide, Volume II](#).

Media type (new media setting)

Specifies the type for the new media that are added to a robot. The list includes the media types that are valid for the robot.

How NetBackup determines the new media type depends on the **Use barcode rules** setting, as follows:

- **Use barcode rules** is selected.
See [“Media type when using barcode rules”](#) on page 483.
- **Use barcode rules** is not selected.
See [“Media type when not using barcode rules”](#) on page 485.

Note: For API robots, the **Media type** is always set to DEFAULT. To specify a media type for API robots, use the **Media Type Mappings** tab of the dialog box.

See [“Configuring media type mappings”](#) on page 496.

Media type when using barcode rules

If you use barcode rules in NetBackup, choose one of the following:

DEFAULT

NetBackup uses the barcode rules to determine the media type that is assigned.

Each media type to be added should have a barcode rule. For example, assume that you want to add DLT and half-inch cartridges to a TLD robot with a single update operation. First create separate barcode rules for DLT and half-inch cartridges and then select the specific media types when you create the barcode rules. Finally, select `DEFAULT` on the **Media Settings** tab. The correct media type is assigned to each media.

If you choose `DEFAULT` on the **Media Settings** tab and `DEFAULT` in the barcode rule, NetBackup assigns the default media type for the robot.

A specific media type from the list.

You can use a single barcode rule to add media of different types, such as DLT and half-inch cartridges (HCART) to a TLD robot. First, select a specific media type on the **Media Settings** tab. Second, select `DEFAULT` for the barcode rule media type when you create the barcode rule. You can perform one update for DLT and another for half-inch cartridge, and the barcode rule assigns the correct media type.

If you specify a value other than `DEFAULT`, the barcode rule media type must be the same as the media or be `DEFAULT`. If not, the barcode rule does not match the media (except for cleaning media).

[Table 9-6](#) shows some combinations of media types on the **Media Settings** tab and barcode rule media types for a TLD (non-API) robot. It also shows the results when the media are added to the volume configuration.

Table 9-6 Example media type and barcode rule combinations

Media type on Media Settings tab	Barcode rule media type	Rule matches?	Media type added to volume configuration
DLT	DEFAULT	Yes	DLT
HCART	DEFAULT	Yes	HCART
DLT	DLT	Yes	DLT
DLT	DLT_CLN	Yes	DLT_CLN
DLT_CLN	DLT	No	DLT_CLN
DLT_CLN	DLT_CLN	Yes	DLT_CLN
DLT_CLN	DEFAULT	Yes	DLT_CLN

Table 9-6 Example media type and barcode rule combinations (*continued*)

Media type on Media Settings tab	Barcode rule media type	Rule matches?	Media type added to volume configuration
DLT	HCART	No	DLT
DEFAULT	DEFAULT	Yes	DLT
DEFAULT	DLT	Yes	DLT
DEFAULT	DLT_CLN	Yes	DLT_CLN
DEFAULT	HCART	No	Depends on robot type

The fourth row in the table shows how both cleaning cartridges and regular volumes are added using one update operation.

All the following conditions must be true:

- The media type on the **Media Settings** tab is for regular media (DLT, in this example).
- The barcode matches a barcode tag.
- The media type for the barcode rule is cleaning media (DLT_CLN).

Another example is available:

The sixth row and seventh row in the table show how to add only a cleaning tape. In the sixth row, you specify the cleaning media type on the **Media Settings** tab and in the barcode rule. In the seventh, specify the cleaning media on the **Media Settings** tab and specify default when you configure the barcode rule.

See [“Configuring barcode rules”](#) on page 490.

Media type when not using barcode rules

Choose one of the following if the barcode rules in NetBackup are not used:

DEFAULT

NetBackup uses the media type that is configured for the drives if:

- The drives in the robot are configured on the robot control host
- All drives the same type
- At least one drive is configured on the robot control host

If the drives are not the same type, NetBackup uses the default media type for the robot.

A specific media type	If the robot supports multiple media types and you do not want to use the default media type, select a specific type. Select a specific media type if: the drives are not configured on the robot control host and the drives are not the default media type for the robot.

The following table shows the default media types for robots when drives are not configured on the robot control host:

Table 9-7 Default media types for non-API robots

Robot type	Default media type
Tape Library DLT (TLD)	DLT cartridge tape. Also supports the following: - DLT cartridge tape 2 and 3, 1/2-inch cartridge tape 1/2-inch cartridge tape 2, 1/2-inch cartridge tape 3

Volume pool (new media setting)

The volume pool for the new media. The actions depend on whether you use barcode rules to assign media attributes, as follows:

DEFAULT	DEFAULT. If you select DEFAULT and: - Use barcode rules, the barcode rules determine the volume pool to which new volumes are assigned - Do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool
A specific volume pool.	If you use barcode rules, this volume pool setting always overrides the rule.

About barcodes

When a robotic library has a barcode reader, it scans the media for barcodes and saves the results. The results associate the slot number and the barcode with the media in that slot. NetBackup obtains the barcode and slot information from the robotic library.

In the robots that have barcode readers, NetBackup verifies the barcode to ensure that the robot loads the correct volume.

If the barcode on the volume does not match the barcode in the NetBackup database, NetBackup does one of the following:

- Assigns the request a pending status (for media-specific jobs such as a restore)
- Uses another volume (for backup or duplicate jobs)

If a requested volume is not in a robot, a pending request message appears in the Device monitor.

The operator must find the volume and do one of the following:

- Check the Device monitor to find a suitable drive and mount the requested volume in that drive.
- Move the volume into the robot, update the volume configuration to reflect the correct location for the media, and resubmit the request.

If the volume is labeled, the automatic volume recognition daemon reads the label and the drive is assigned to the request. If the volume is unlabeled and not associated with a robot, the operator manually assigns the drive to the request.

Barcode advantages

NetBackup functions well whether or not barcodes are used. However, it is recommended to use a media with barcodes in the robots that can read barcodes.

Barcodes offer the following advantages:

- Automatic media ID assignment
When you add new media to a robot, NetBackup is able to assign media IDs according to specified criteria.
- More accurate tracking of volume location
A robot inventory update can determine which volumes are in a robot.
- Increased performance
Not using barcodes can adversely affect performance for some robots. A robot that reads barcodes performs a scan each time it moves a tape. The robot stores the correct barcode in memory or verifies a previously saved barcode. However, if a tape does not have a barcode, the robot retries the scan multiple times, degrading performance.

Barcode best practices

Consider the following practices when you select barcodes for volumes:

- Barcodes usually appear on the labels that are attached to the outside of tape volumes.
- The maximum barcode length that NetBackup supports depends on the type of robot.

For more information, see the [NetBackup Device Configuration Guide](#).

- Always follow the robotic library vendor's recommendations when purchasing barcode labels for use with NetBackup.
Ensure that the barcodes have the correct number of characters.
- Barcodes can represent any combination of alpha and numeric characters, but different robots support different lengths of barcodes.
See the robot vendor's documentation to determine the requirements for a specific robot type.
- Use barcodes without spaces (at the beginning, at the end, or between any characters).
Otherwise, the robot or NetBackup may not read them correctly.
- Volumes in an API robot have a real or a logical barcode.
This volume identifier is used as the NetBackup media ID. This volume identifier is the volume serial number in ACS robots.
- For API robots, the barcode for a volume must be identical to the NetBackup media ID.
Match barcodes to media IDs by getting custom labels in the same series as the media IDs. For example, to match a set of media IDs from AA0000 to ZZ9999, get barcode labels in that series.
- When a robotic library can contain more than one media type, assign specific characters in the barcode to different media types. Do so by using media ID generation rules.
Also, use barcodes to differentiate between data tapes and cleaning tapes or to differentiate between volume pools.

About barcode rules

A barcode rule specifies criteria for assigning attributes to new robotic volumes. NetBackup assigns these attributes by using the barcode for the volume that the robotic library provides and your barcode rules.

In NetBackup, you choose whether to use barcode rules when you set up the robot inventory update operation. The barcode rules are stored on the primary server.

Note: NetBackup does not use barcode rules if a volume already uses a barcode.

About NetBackup actions for barcodes

When a robot inventory update operation uses NetBackup barcode rules and a new barcode is detected in the robot, NetBackup does the following:

- Searches the list of rules (from first to last) for a rule that matches the new barcode.
- If the barcode matches a rule, NetBackup verifies that the media type in the rule is compatible with the media type specified for the update.
- If the media types match, NetBackup assigns the attributes in the rule to the volume. The attributes include the media type, volume pool, maximum number of mounts (or number of cleanings), and description.

Example barcode rules

The following table shows some example barcode rules. Rules are sorted first according to the number of characters in the barcode tag and then by the order added. Two exceptions are the <NONE> and <DEFAULT> rules, which are always located at the end of the list.

Table 9-8 Example barcode rules

Barcode tag	Media type	Volume pool	Max mounts and cleanings	Description
DLT	DLT	d_pool	200	DLT backup
CLD	DLT_CLN	None	30	DLT cleaning
<NONE>	DEFAULT	None	0	No barcode
<DEFAULT>	DEFAULT	NetBackup	0	Other barcodes

Assume that you select the following media settings (update options) for the update operation for a new HCART volume in a TLD robot:

Media type = HCART

Volume group = 00_000_TLD

Use barcode rules = YES

Volume pool = DEFAULT

If a new volume in this robotic library has a barcode of TLD00001, NetBackup uses the rule with the barcode tag of TLD. NetBackup assigns the following attributes to the volume:

- Media ID = 800001 (last six characters of barcode)
- Volume group = 00_000_TLD
- Volume pool = t_pool
- Maximum mounts = 0 (no maximum)

If a new volume has a barcode of TL000001, NetBackup uses the rule with the barcode tag of TL. NetBackup assigns the following attributes to the volume:

- Media ID = 000001 (last six characters of barcode)
- Volume group = 00_000_TLD
- Volume pool = None
- Maximum mounts = 0 (no maximum)

Configuring barcode rules

Use the **Barcode rules** settings in the Inventory robot **Advanced options** to configure the rules that assign attributes to the new volumes that are added to a robot. NetBackup assigns barcodes when you select **Use barcode rules** in the **Media settings**.

To enable barcode rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 418.

For more information about the `vm.conf` file, see the [NetBackup Administrator's Guide, Volume II](#).

See [“About barcodes”](#) on page 486.

To configure barcode rules

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**.
- 3 On the top right, select **Actions > Inventory robots**.
- 4 Select the **Device host**.
- 5 Select the robot.
- 6 Select either **Preview volume configuration changes** or **Update volume configuration**.
- 7 Select **Advanced options** to open the menu. Then select **Barcode rules**.

- 8
- To configure the rules, do any of the following:
- Add a rule

Select **Add** and then configure the rule.
- Edit a rule

Locate the rule and select **Edit**, then change the rule.
- Delete a rule

Locate the rule, select **Delete**.
- See “Barcode rules settings” on page 491.
- 9
- When you are finished configuring rules, select **Save**.

Barcode rules settings

The following table describes the settings you can configure for barcode rules. NetBackup uses these rules to assign barcodes to new media.

Table 9-9 Barcode rule settings

Barcode rule setting	Description
Barcode tag	A unique string of barcode characters that identifies the type of media. For example, use DLT as the barcode tag for a barcode rule if the following is true: - You use DLT on the barcodes to identify DLT tapes - DLT is not used on any other barcodes in the robot Similarly, if you use CLND for DLT cleaning media, use CLND as the barcode tag for the rule for DLT cleaning media. The barcode tag can have from 1 to 16 characters but cannot contain spaces. The following are the special barcode rules that can match special characters in the barcode tags: - NONE Matches when rules are used and the volume has an unreadable barcode or the robot does not support barcodes. - DEFAULT For volumes with barcodes, this tag matches when none of the other barcode tags match. However, the following must be compatible: the media type in the DEFAULT rule and the media type on the Media Settings tab. You cannot change the barcode tag of a barcode rule. Instead, first delete the old rule, then add a rule with a new barcode tag. Use the Media Settings tab to set up the criteria for a robot update. See “Configuring media settings” on page 478.
Description	A description of the barcode rule. Enter from 1 to 25 characters.

Table 9-9 Barcode rule settings (*continued*)

Barcode rule setting	Description
Maximum mounts	The maximum number of mounts (or cleanings) that are allowed for the volume. For data volumes, a value of zero means the volume can be mounted an unlimited number of times. For cleaning tapes, zero means that the cleaning tape is not used. It is recommended that you use barcodes for the cleaning media that cannot be confused with barcodes for data media. Doing so can avoid a value of 0 for cleaning tapes.
Media type option	The media type to assign to the media. The media type that is specified on the Media Settings tab always overrides the media type of the barcode rule. If you specify a value other than <code>DEFAULT</code> on the Media Settings tab, the barcode rule media type must be the same as the media or be <code>DEFAULT</code>. If not, the barcode rule does not match the media (except for cleaning media). See “Media type when using barcode rules” on page 483. Note: When a media type is selected, the maximum mounts value may revert to the default value for the specified media type. For example, it may revert to 0 for unlimited when you select a non-cleaning media type. See “NetBackup media types” on page 466.
Volume pool	The volume pool for the new media. The actions depend on whether you use barcode rules to assign media attributes. Select from the following: ■ DEFAULT If DEFAULT is selected, NetBackup performs the following actions: ■ If you use barcode rules, the barcode rules determine the volume pool to which new volumes are assigned. ■ If you do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool. ■ A specific volume pool This volume pool setting always overrides any barcode rules.

About media ID generation rules

Use media ID generation rules to override the default media ID naming method NetBackup uses. The default method uses the last six characters of the barcode the robot provides to generate the media ID.

Note: To use media ID generation rules, the robot must support barcodes and the robot cannot be an API robot. Media ID generation rules are saved in the Media Manager configuration file (`vm.conf`). Information about the `vm.conf` file is in the [NetBackup Administrator's Guide, Volume II](#).

For example, two eight-character barcodes are `S000006L11` and `0000006L11`. Without any media ID generation rules NetBackup uses the last six characters of the barcode to generate media IDs. In this example, the same media ID for the two barcodes is created (`0006L11`).

Use a rule to control how NetBackup creates media IDs by specifying which characters of a barcode are used in the media ID. Or, specify that alphanumeric characters are to be inserted into the ID.

Define multiple rules to accommodate the robots and the barcode lengths. Define rules to specific robots and for each barcode format that has different numbers or characters in the barcode. Multiple rules allow flexibility for the robots that support multiple media types.

Configuring media ID generation rules

For non-API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 418.

Use the **Media ID Generation** tab of the NetBackup **Advanced Robot Inventory Options** dialog box to configure the rules that override the default naming method. To use media ID generation rules, the robot must support barcodes and the robot cannot be an API robot.

See [“About media ID generation rules”](#) on page 492.

To configure media ID generation rules

1 Open the **Advanced Robot Inventory Options** dialog box, as follows:

- From the **Robot Inventory** dialog box
- 1**

In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots** in the left pane.

2

Select the robot that you want to inventory.

3

On the **Actions** menu, select **Inventory Robot**.

4

Click either **Preview volume configuration changes** or **Update volume configuration**.

5

Click **Advanced Options**.


- From the **Volume Configuration Wizard**
- 1

In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Devices**.
- 2

From the list of wizards in the right pane, click **Configure Volumes** and follow the wizard instructions.
- 3

On the **Robot Inventory** panel of the **Volume Configuration Wizard**, click **Advanced Options**.

- 2
- In the **Advanced Robot Inventory Options** dialog box, click the **Media ID Generation** tab.



3 To configure the rules, do any of the following:

Add a rule	Click New and then configure the rule in the dialog box.
Change a rule	Select the rule, click Change , and then change the rule in the dialog box. You cannot change the robot number or barcode length of a rule. To change those properties, first delete the old rule and then add a rule.
Delete a rule	Select the rule, click Delete , and click OK in the confirmation dialog box. You can select and delete multiple rules with one operation.

See [“Media ID generation options”](#) on page 495.

4 When you are finished configuring rules, click **OK**.

Media ID generation options

NetBackup uses rules to generate the IDs for media in robots. The default rule uses the last six characters of the barcode label from the tape.

You can configure media ID generation rules to override the default rule. Control how NetBackup creates media IDs by defining the rules that specify which characters of a barcode label to use for the media ID.

The following subsections describe the media ID generation rule options.

The following list describes the media ID generation rule options:

- **Bar code length**
The **Barcode length** is the number of characters in the barcode for tapes in the robot.
You cannot change the barcode length of a rule. Rather, first delete the rule and then add a new rule.
- **Media ID generation rule**
A **Media ID generation rule** consists of a maximum of six colon-separated fields. Numbers define the positions of the characters in the barcode that are to be extracted. For example, the number 2 in a field extracts the second character (from the left) of the barcode. You can specify numbers in any order.
To insert a specific character in a generated media idea, precede the character by a pound sign (#). Any alphanumeric characters that are specified must be valid for a media ID.
Use rules to create media IDs of many formats. However, it may be difficult to manage media if the label on the media and the generated media ID are different. The table shows some examples of rules and the resulting media IDs.

Barcode on tape	Media ID generation rule	Generated media ID
032945L1	1:2:3:4:5:6	032945
032945L1	3:4:5:6:7	2945L
032945L1	#N:2:3:4:5:6	N32945
543106L1	#9:2:3:4	9431
543106L1	1:2:3:4:#P	5431P

- **Robot number**

The number of the robot to which the rule applies.

You cannot change the robot number of a rule. Rather, first delete the rule and then add a new rule.

About media type mapping rules

Applies to API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 418.

For API robots, NetBackup contains default mappings from a vendor's media types to NetBackup media types. API robots are ACS robot types.

You can change the default mappings. Changes apply only to the current volume configuration update.

You also can add media type mappings.

See [“About adding media type mapping entries”](#) on page 499.

See [“Default and allowable media types”](#) on page 499.

See [“NetBackup media types”](#) on page 466.

Note: You can write a barcode rule that contains the media types that are incompatible with vendor media types. However, the robot inventory update may assign NetBackup media types that are inconsistent with the vendor media types. Avoid this problem by grouping barcode rules by media type.

Configuring media type mappings

Applies to API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 418.

Use the **Media Type Mappings** tab of the NetBackup **Advanced Robot Inventory Options** dialog box to configure the attributes for existing and new media.

See [“About media type mapping rules”](#) on page 496.


To configure media type mappings

- 1 Open the **Advanced Robot Inventory Options** dialog box, as follows:

From the **Robot Inventory** dialog box

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots** in the left pane.
- 2 Select the robot that you want to inventory.
- 3 On the **Actions** menu, select **Inventory Robot**.
- 4 Click either **Preview volume configuration changes** or **Update volume configuration**.
- 5 Click **Advanced Options**.

- From the **Volume Configuration Wizard**
 - 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Devices**.
 - 2 From the list of wizards in the right pane, click **Configure Volumes** and follow the wizard instructions.
 - 3 On the **Robot Inventory** panel of the **Volume Configuration Wizard**, click **Advanced Options**.
- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Media Type Mappings** tab.



The mappings that appear are only for the robot type that was selected for inventory. The default mappings and any mappings you added or changed appear.

- 3 Select the row that contains the robot-vendor media type mapping that you want to change and click **Change Mapping**.
- 4 In the **Change Media Mapping** dialog box, select a media type from the list of allowed selections.
- 5 Click **OK**.

To reset the mappings to the default, click **Reset to Defaults**.

About adding media type mapping entries

Applies to API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 418.

The default media type mappings may not provide the wanted mappings. If not, add robot-specific media mappings to the `vm.conf` file on the host on which you run the NetBackup web UI.

For more information about the `vm.conf` file, see the [NetBackup Administrator's Guide, Volume II](#).

Table 9-10 Examples of robot-specific media mappings

vm.conf entry	Result	Robot default without a vm.conf entry
ACS_3490E = HCART2	Maps the ACS 3490E to the HCART2 media type.	HCART
ACS_DLTIV = DLT2	Maps ACS DLTIV to the DLT2 media type.	DLT for all ACS DLT media types, including DLTIV

Default and allowable media types

Applies to API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 418.

The default media types are provided by each robot vendor.

[Table 9-11](#) contains the default and allowable media types for the API robots.

The following items provide information to help you understand the table:

- The first column of the table shows the vendor's media type.
- The second column of the table shows the default media type in NetBackup.
- The third column shows the media types to which you can map the defaults. To do so, first add the allowable mapping entries to the `vm.conf` file.

Some map entries are not allowed. For example, you cannot specify either of the following map entries for ACS robots:

```
ACS_DD3A = DLT
ACS_DD3A = HCART4
```

Table 9-11 Default and allowable media types for ACS robots

ACS media type	Default media type	Allowable media types through mappings
3480	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3490E	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
DD3A	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
DD3B	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
DD3C	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
DD3D	1/2-inch cartridge cleaning tape 2 (HC2_CLN)	HC_CLN, HC2_CLN, HC3_CLN
DLTIII	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
DLTIIIXT	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
DLTIV	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
EECART	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
JLABEL	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
KLABEL	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_100G	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_10GB	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_200G	1/2-inch cartridge (HCART2)	HCART, HCART2, HCART3
LTO_35GB	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3

Table 9-11 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
LTO_400G	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
LTO_400W	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
LTO_50GB	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_800G	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_800W	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_1_5T	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
LTO_1_5W	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
LTO_2_5T	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
LTO_2_5W	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
LTO_6_4T	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_6_4W	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_CLN1	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
LTO_CLN2	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
LTO_CLN3	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
LTO_CLNU	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
SDLT	Digital Linear Tape 3 (DLT3)	DLT, DLT2, DLT3

Table 9-11 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
SDLT_2	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_4	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_S1	Digital Linear Tape 2 (DLT2)	DLT, DLT2, DLT3
SDLT_S2	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_S3	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_S4	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
STK1R	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
STK1U	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
STK1Y	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
STK2P	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
STK2W	1/2-inch cartridge cleaning tape 2 (HC2_CLN)	HC_CLN, HC2_CLN, HC3_CLN
T10000CC	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T10000CL	1/2-inch cartridge cleaning tape 3 (HC3_CLN)	HC_CLN, HC2_CLN, HC3_CLN
T10000CT	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T10000T1	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T10000T2	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T10000TS	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3

Table 9-11 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
T1000TT	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
UNKNOWN (for unknown ACS media types)	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3, DLT, DLT2, DLT3
VCART	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
VIRTUAL	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3

Managing volumes

The following sections describe the procedures to manage volumes.

Changing the group of a volume

If you move a volume physically to a different robot, change the group of the volume to reflect the move.

See [“About rules for moving volumes between groups”](#) on page 504.

To change the group of a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select the volumes that you want to change the volume group assignment for.
- 5 Select **Change volume group**.
- 6 For the **Volume group**, enter the name of the new volume group. Or, select a name from the list.
- 7 Select **Confirm**.

The name change is reflected in the volume list entry for the selected volumes. If you specified a new volume group (which creates a new volume group), the group appears under **Volume groups**.

About rules for moving volumes between groups

The following are the rules for moving volumes between groups:

- The target volume group must contain the same type of media as the source volume group. If the target volume group is empty: The successive volumes that you add to it must match the type of media that you first add to it.
- All volumes in a robotic library must belong to a volume group. If you do not specify a group, NetBackup generates a new volume group name by using the robot number and type.
- More than one volume group can share the same location. For example, a robotic library can contain volumes from more than one volume group and you can have more than one standalone volume group.
- All members of a group must be in the same robotic library or be standalone. That is, if volume group already exists in another robotic library, you cannot add it (or part of it) to a robotic library.

See [“About NetBackup volume groups”](#) on page 466.

See [“About moving volumes”](#) on page 517.

Changing the owner of a volume

You can change the media server or server group that owns the volume.

See [“About NetBackup server groups”](#) on page 374.

See [“About media sharing”](#) on page 525.

To change the owner of a volume

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the **Volumes** list, select the volume that you want to change.
- 3 On the **Actions** menu, select **Change Media Owner**.

4 In the **Media Owner** field, select one of the following:

Any (default)	Allows NetBackup to choose the media owner. NetBackup chooses a media server or a server group (if one is configured).
None	Specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	Specify a server group. A server group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the drop-down list.

5 Click **OK**.

Edit a volume

You can change some of the properties of a volume, including the volume pool.

To change volume properties

- 1** Open the **NetBackup web UI**.
- 2** On the left, select **Storage > Tape storage**.
- 3** Select the **Volumes** tab.
- 4** Locate and select the volume. Select **Edit**.
- 5** Change the properties for the volume.
See [“Volume properties”](#) on page 475.
- 6** Select **Update**.

About assigning and deassigning volumes

An assigned volume is one that is reserved for exclusive use by NetBackup. A volume is set to the assigned state when either application writes data on it for the first time. The time of the assignment appears in the **Time assigned** column for the volume on the **Volumes** tab. When a volume is assigned, you cannot delete it or change its volume pool.

A volume remains assigned until NetBackup deassigns it.

NetBackup deassigns a volume only when the data is no longer required, as follows:

- For regular backup volumes, when the retention period has expired for all the backups on the volume.
- For catalog backup volumes, when you stop using the volume for catalog backups.

To deassign a volume, you expire the images on the volume. After you expire a volume, NetBackup deassigns it and does not track the backups that are on it. NetBackup can reuse the volume, you can delete it, or you can change its volume pool.

See [“Expiring backup images”](#) on page 933.

You can expire backup images regardless of the volume state (Frozen, Suspended, and so on).

NetBackup does not erase images on expired volumes. You can still use the data on the volume by importing the images into NetBackup (if the volume has not been overwritten).

See [“About importing backup images”](#) on page 935.

Note: It is not recommended that you deassign NetBackup volumes. If you do, be certain that the volumes do not contain any important data. If you are uncertain, copy the images to another volume before you deassign the volume.

Delete a volume

You can delete volumes from the NetBackup configuration. For example, if any of the following situations apply, you may want to delete the volume:

- A volume is no longer used and you want to recycle it by relabeling it with a different media ID.
- A volume is unusable because of repeated media errors.
- A volume is past its expiration date or has too many mounts, and you want to replace it with a new volume.
- A volume is lost and you want to remove it from the NetBackup database.

After a volume is deleted, you can discard it or add it back under the same or a different media ID.

See [“About assigning and deassigning volumes”](#) on page 505.

To delete volumes

- 1 Before you delete and reuse or discard a volume, ensure that it does not have any important data. You cannot delete NetBackup volumes if they are assigned.
- 2 Open the **NetBackup web UI**.
- 3 Click **Storage > Tape storage**.
- 4 Click the **Volumes** tab.
- 5 Select the desired volume from the volumes list and click **Delete > Delete**.
- 6 Remove the deleted volume or volumes from the storage device.

Erasing a volume

You can erase the data on a volume if the following are true:

- The volume is not assigned.
- The volume contains no valid NetBackup images.

After NetBackup erases the media, NetBackup writes a label on the media.

If you erase media, NetBackup cannot restore or import the data on the media.

If a volume contains valid NetBackup images, deassign the volume so NetBackup can label it.

See [“About assigning and deassigning volumes”](#) on page 505.

The following table describes the types of erase.

Table 9-12 Types of erase

Type of erase	Description
SCSI long erase	Rewinds the media and the data is overwritten with a known data pattern. A SCSI long erase is also called a secure erase because it erases the recorded data completely. Note: A long erase is a time-consuming operation and can take as long as two hours to three hours. For example, it takes about 45 minutes to erase a 4-mm tape on a standalone drive
SCSI quick erase	Rewinds the media and an erase gap is recorded on the media. The format of this gap is drive dependent. It can be an end-of-data (EOD) mark or a recorded pattern that the drive does not recognize as data. Some drives do not support a quick erase (such as QUANTUM DLT7000). For the drives that do not support a quick erase, the new tape header that is written acts as an application-specific quick erase.

Note: NetBackup does not support erase functions on NDMP drives.

To erase a volume

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, select a volume or volumes that you want to erase.

If you select multiple volumes, they must all be in the same robot.

- 3 Select either **Actions > Quick Erase** or **Actions > Long Erase**.

- 4 In the erase dialog box, specify the name of the media server to initiate the erase operation.

To overwrite any existing labels on the media, do not select **Verify media label before performing operation**.

- 5 Click **OK**.

A dialog box warns you that this action is irreversible.

- 6 Click **OK** if you are certain you want to start the erase action.

A dialog box reminds you to use the **Activity Monitor** to view the progress and status of the action. (For many types of drives, you may not be able to cancel a label or erase media job from the **Activity Monitor**.) Click **OK**.

If you selected **Verify media label before performing operation** and the actual volume label does not match the expected label, the media is not erased.

About exchanging a volume

You should exchange a volume (replace one volume with another volume) if a volume meets any of the following conditions:

- Full (in this case, to exchange a volume means to remove the volume from a robotic tape library).
- Past the maximum number of mounts.
- Old (past the expiration date).
- Unusable (for example, because of repeated media errors).

Depending on whether you want to reuse the old media ID or not, follow one of the exchange volumes processes in the following subsections.

Exchanging a volume and using a new media ID

Use this procedure when the following are true:

- The volume contains current and valid NetBackup images.
- You require slots in the robotic library for additional backups, duplications, vault functions, or other purposes.

The following table describes the procedure that used to exchange a volume and use a new media ID.

Table 9-13 Exchange a volume and using a new media ID

Step	Task	Instructions
Step 1	Move the volume to another location If the volume is in a robotic library, remove it from the robotic library and move it to a standalone group.	See “About moving volumes” on page 517.
Step 2	Add a new volume or move an existing volume in as a replacement for the volume you removed. If you add a new volume, specify a new media ID. Specify the same values for the other attributes as the removed volume (such as robotic residence, volume pool, and the media type).	See “About adding volumes” on page 472.
Step 3	Physically replace the old volume. Do not delete the old volume in case you need to retrieve the data on the volume.	Beyond the scope of the NetBackup documentation.

Exchanging a volume and using the old media ID

You can exchange a volume and reuse the same media ID, which may be convenient in some instances.

Reuse a media ID only if all data on the old volume is not required and you recycle or discard the volume.

Warning: If you exchange a media ID for a volume that has unexpired backup images, serious operational problems and data loss may occur.

The following table describes the procedure to exchange a volume and use the old media ID.

Table 9-14 Exchange a volume and use the old media ID

Step	Task	Instructions
Step 1	Delete the volume.	See “Delete a volume” on page 506.
Step 2	Remove the old volume from the storage device. Physically add the new volume to the storage device.	See “About injecting and ejecting volumes” on page 511.
Step 3	Add the new volume to the NetBackup volume configuration and specify the same attributes as the old volume, including the old media ID.	See “About adding volumes” on page 472.
Step 4	Set a new expiration date for the volume.	See “Edit a volume” on page 505.
Step 5	Optionally, label the volume. Although you do not have to label the volume, the label process puts the media in a known state. The external media label matches the recorded media label, and the mode is known to be compatible with the drives in the robotic library.	See “Labeling a volume” on page 516.

About frozen media

Frozen media is the media that NetBackup does not use for backups. NetBackup stops directing the backups and the archives to frozen media. NetBackup never deletes a frozen media ID from the NetBackup media catalog, even after the retention period ends for all backups on the media. NetBackup does not unassign a frozen volume from the NetBackup volume pool when its backup images expire.

All unexpired backup images on frozen media continue to be available for restores.

NetBackup freezes the tape volumes for a variety of reasons, as follows:

- NetBackup freezes a volume when read or write errors surpass the threshold within the time window. The default media error threshold is 2. That is, NetBackup freezes media on the third media error in the default time window (12 hours). Common reasons for write failures are dirty write heads or old media. The reason for the action is logged in the NetBackup error catalog (view the Media Logs report or the All Log Entries report).

You can use the NetBackup `nbemmcmd` command with the `-media_error_threshold` and `-time_window` options to change the default values.

For more information about the `nbemmcmd` command, see the [NetBackup Commands Reference Guide](#).

- NetBackup freezes a volume if a write failure makes future attempts at positioning the tape unreliable.

- NetBackup freezes the catalog volumes during catalog recovery.
- NetBackup freezes the volumes in some circumstances with write once read many (WORM) media or WORM-capable drives.
See [“About using volume pools to manage WORM media”](#) on page 469.

You can unfreeze the frozen volumes manually.

See [“Freezing or unfreezing a volume”](#) on page 511.

Freezing or unfreezing a volume

NetBackup freezes volumes under circumstances.

Use the following procedure to manually freeze or unfreeze a volume.

To freeze or unfreeze media

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, select the volume that you want to freeze or unfreeze.
- 3 On the **Actions** menu, select **Freeze** or **Unfreeze**.
- 4 In the dialog box, click **OK**.

About injecting and ejecting volumes

Media access port (MAP) functionality differs between robotic libraries. For many libraries, NetBackup opens and closes the MAP as needed. However, some libraries have the front-panel inject and the eject functions that conflict with NetBackup's use of the media access port. And for other libraries, NetBackup requires front-panel interaction by an operator to use the media access port.

Read the operator manual for the library to understand the media access port functionality. Some libraries may not be fully compatible with the inject and eject features of NetBackup unless properly handled. Other libraries may not be compatible at all.

Injecting volumes into robots

You can inject volumes into the robots that contain media access ports.

Any volumes to be injected must be in the media access port before the operation begins. If no volumes are in the port, you are not prompted to place volumes in the media access port and the update operation continues.

Each volume in the MAP is moved into the robotic library. If the MAP contains multiple volumes, they are moved to empty slots in the robotic library until the media access port is empty or all the slots are full.

After the volume or volumes are moved, NetBackup updates the volume configuration.

Some robots report only that media access ports are possible. Therefore, the **Empty media access port prior to update** option may be available for some robots that do not contain media access ports.

To inject volumes into the robots that contain media access ports

- 1 Load the volumes in the MAP.
- 2 Inventory the robot.

See [“Updating the NetBackup volume configuration with a robot’s contents”](#) on page 538.

- 3 Select **Empty media access port prior to update**.

Ejecting volumes

You can eject single or multiple volumes.

You cannot eject multiple volumes with one operation if they reside in multiple robots.

Operator intervention is only required if the robotic library does not contain a media access port large enough to eject all of the selected volumes. For these robot types, NetBackup prompts an operator to remove the media from the media access port so the eject operation can continue.

See [“Media ejection timeout periods”](#) on page 513.

To eject volumes

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volumes** tab.
- 3 Select one or more volumes that you want to eject.

4 Click **Eject from robot**.

5 Do one of the following actions:

ACS robots	Select the media access port to use for the ejection, then click Eject .
TLD robots	Click Eject .

The robotic library may not contain a media access port large enough to eject all of the selected volumes. For most robot types, you are prompted to remove the media from the media access port so the eject can continue with the remaining volumes.

See “[NetBackup robot types](#)” on page 418.

Media ejection timeout periods

The media ejection period (the amount of time before an error condition occurs) varies depending on the capability of each robot.

The following table shows the ejection timeout periods for robots.

Table 9-15 Media ejection timeout periods

Robot types	Timeout period
Automated Cartridge System (ACS)	One week
Tape Library DLT (TLD)	30 minutes.

Note: If the media is not removed and a timeout condition occurs, the media is returned to (injected into) the robot. Inventory the robot and eject the media that was returned to the robot.

Some robots do not contain media access ports. For these robots, the operator must remove the volumes from the robot manually.

Note: After you add or remove media manually, use NetBackup to inventory the robot.

About rescanning and updating barcodes

You can rescan the media in a robot and then update NetBackup with the barcodes of that media.

You should rescan and update only in certain circumstances.

Note: Rescan and update barcodes does not apply to volumes in API robot types.

See [“NetBackup robot types”](#) on page 418.

When to rescan and to update barcodes	Rescan and update barcodes only to add the barcodes that are not in the NetBackup database. For example: if you add a new volume but do not insert the tape into the robot, NetBackup does not add the barcode to the database. Use this command to add the barcode after you insert the tape into the robotic library.
When not to rescan and to update barcodes	Do not rescan and update to correct the reports that show a media ID in the wrong slot. To correct that problem, perform one of the following actions: ■ Logically move the volume by selecting a volume and then on the Actions menu select Move. ■ Logically move the volume by updating the volume configuration. See “Updating the NetBackup volume configuration with a robot's contents” on page 538. ■ Physically move the volume into the correct slot. To obtain an inventory of the robot without updating the barcode information in the database, inventory the robot and use the show contents option. See “Showing the media in a robot” on page 533.

Rescanning and updating barcodes

Use the following procedure to rescan the media in a robot and to update NetBackup with the barcodes.

Note: Rescan and update barcodes does not apply to volumes in API robot types.

See [“NetBackup robot types”](#) on page 418.

To rescan barcodes and update the NetBackup database

- 1** In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media > Robots**.
- 2** Select the robotic library that contains the volumes that you want to scan and update.
- 3** In the right pane, in the **Volumes** list, select the volumes.
- 4** On the **Actions** menu, select **Rescan/Update Barcodes**.
- 5** Click **Start**.

The results of the update appear in the output section of the dialog box.

About labeling NetBackup volumes

When NetBackup labels a volume, it writes a record on the magnetic tape of the volume; the record (or label) includes the NetBackup media ID.

Normally, NetBackup controls the labeling of its volumes. In most cases, NetBackup labels a volume the first time it is used for a backup.

The volume label depends on whether or not the media has a barcode, as follows:

- If the robot supports barcodes and the media has barcodes, NetBackup uses the last six characters of the barcode for the media ID.
To change this default action, specify and select specific characters by using Media ID generation rules.
See [“Configuring media ID generation rules”](#) on page 493.
- For volumes without barcodes, by default NetBackup uses a prefix of the letter A when it assigns a media ID to a volume (for example, A00001).
To change the default prefix, use the `MEDIA_ID_PREFIX` configuration option in the `vm.conf` file.
For more information about the `vm.conf` file, see the [NetBackup Administrator's Guide, Volume II](#).

Media is not labeled automatically in the following situations:

- They were last used for NetBackup catalog backups.
Do not label catalog backup volumes unless they are no longer used for catalog backups.
- They contain data from a recognized non-NetBackup application and NetBackup is configured to prohibit media overwrite for that media type.

To label these media types, the following must be true:

- NetBackup has not assigned the media

- The media contains no valid NetBackup images

Labeling a volume

If a volume contains valid NetBackup images, deassign the volume so that it can be labeled.

See [“About assigning and deassigning volumes”](#) on page 505.

If you want to label media and assign specific media IDs (rather than allow NetBackup to assign IDs), use the `bplabel` command.

Note: If you label a volume, NetBackup cannot restore or import the data that was on the media after you label it.

Note: For many types of drives, you may not be able to cancel a label job from the Activity Monitor.

See [“About labeling NetBackup volumes”](#) on page 515.

To label a volume

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, select a volume or the volumes that you want to label.

If you select multiple volumes, they all must be in the same robot.
- 3 On the **Actions** menu, select **Label**.
- 4 In the **Label** dialog box, specify the following properties for the label operation.

Media server	Enter tname of the media server that controls the drive to write the label.
Verify label before performing operation	Select this option to verify that the media in the drive is the expected media. To overwrite any existing labels on the media, do not select Verify media label before performing operation .

- 5 Click **OK**.
- 6 In the warning dialog box, click **OK**.

If you selected **Verify media label before performing operation** and the actual volume label does not match the expected label, the media is not relabeled.

About moving volumes

When you move volumes in or out of a robotic library or from one robot to another, move the volumes physically and logically, as follows:

- Physically move volumes by inserting or by removing them. For some robot types, use the NetBackup inject and eject options.
- Logically move volumes using NetBackup, which updates the NetBackup database to show the volume at the new location.

When you move volumes from one robotic library to another robotic library, perform the following actions:

- Move the volumes to stand alone as an intermediate step.
- Move the volumes to the new robotic library.

The following types of logical moves are available:

- Move single volumes.
- Move multiple volumes.
- Move combinations of single and multiple volumes.
- Move volume groups.

You cannot move volumes to an invalid location.

It is recommended that you perform moves by selecting and by moving only one type of media at a time to a single destination.

The following are several examples of when to move volumes logically:

- When a volume is full in a robotic library and no slots are available for new volumes in the robotic library. Move the full volume to stand alone, remove it from the robot, then configure a new volume for the empty slot or move an existing volume into that slot. Use the same process to replace a defective volume.
- Moving volumes from a robotic library to an off-site location or from an off-site location into a robotic library. When you move tapes to an off-site location, move them to stand alone.

- Moving volumes from one robotic library to another (for example, if a library is down).
- Changing the volume group for a volume or volumes.

Moving volumes by using the robot inventory update option

Use this procedure for the following:

To move volumes within a robot.	The robot must have a barcode reader and the volumes must contain readable barcodes.
To remove volumes from a robot.	Use this procedure even if the volumes do not contain barcodes or if the robot does not have a reader.

To move volumes by using a robot inventory update

- 1 Physically move the volumes to their new location.
- 2 On the **Actions** menu, select **Inventory Robot**.
- 3 In the **Robot Inventory** dialog box, select **Update volume configuration**.
- 4 Select other options as appropriate.

See [“About robot inventory”](#) on page 528.

Moving volumes by using the Actions menu

If you move a volume to a robotic library that has a barcode reader, NetBackup updates the NetBackup database with the correct barcode.

To move volumes by using the Actions menu

- 1 Physically move the volumes to their new location.
- 2 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 3 In the right pane, in the **Volumes** list, select the volumes that you want to move.
- 4 On the **Actions** menu, select **Move**.
- 5 In the **Move Volumes** dialog box, specify the properties for the move.

If you move a single volume, the dialog box entries show the current location of the volume.

See [“Volume properties”](#) on page 475.

About recycling a volume

If you recycle a volume, you can use either the existing media ID or a new media ID.

Caution: Recycle a volume only if all NetBackup data on the volume is no longer needed or if the volume is damaged and unusable. Otherwise, you may encounter serious operational problems and a possible loss of data.

Recycling a volume and using the existing media ID

NetBackup recycles a volume and returns it to the volume rotation when the last valid image on the volume expires.

To recycle a volume that contains unexpired backup images, you must deassign the volume.

See [“About assigning and deassigning volumes”](#) on page 505.

Recycling a volume and using a new media ID

Recycle a volume if it is a duplicate of another volume with the same media ID. Also recycle a volume if you change how you name volumes and you want to match the barcodes on the volume.

The following table describes the procedure to recycle a volume and use a new media ID.

Table 9-16 Recycling a volume and using a new media ID

Step	Action	Description
Step 1	Physically remove the volume from the storage device.	See “Ejecting volumes” on page 512.
Step 2	If the volume is in a robotic library, move it to standalone.	See “About moving volumes” on page 517.
Step 3	Record the current number of mounts and expiration date for the volume.	Go to Storage > Tape storage > Volumes in the NetBackup web UI .
Step 4	Delete the volume entry.	See “Delete a volume” on page 506.

Table 9-16 Recycling a volume and using a new media ID (*continued*)

Step	Action	Description
Step 5	Add a new volume entry.	See “Adding volumes” on page 474. Because NetBackup sets the mount value to zero for new volume entries, you must adjust the value to account for previous mounts. Set the maximum mounts to a value that is equal to or less than the following value: The number of mounts that the manufacturer recommends minus the value that you recorded earlier.
Step 6	Physically add the volume to the storage device.	See “Injecting volumes into robots” on page 511.
Step 7	Configure the number of mounts.	Set the number of mounts to the value you recorded earlier by using the following command: On Windows hosts: <pre>install_path\Volmgr\bin\vmchange -m media_id -n number_of_mounts</pre> On UNIX hosts: <pre>/usr/openv/volmgr/bin/vmchange -m media_id -n number_of_mounts</pre>
Step 8	Set the expiration date to the value you recorded earlier.	See “Edit a volume” on page 505.

Suspending or unsuspending volumes

You cannot use a suspended volume for backups until retention periods for all backups on it have expired. At that time, NetBackup deletes the suspended volume from the NetBackup media catalog and unassigns it from NetBackup.

A suspended volume is available for restores. If the backups have expired, import the backups first.

To suspend or unsuspend media

- 1 In the **NetBackup Administration Console**, in the left pane, select **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, select the volume or volumes that you want to suspend or unsuspend.

- 3 On the **Actions** menu, select **Suspend** or **Unsuspend**.
- 4 In the dialog box, click **OK**.

Managing volume pools

The following sections describe the operations you can perform to manage volume pools.

See [“Add a volume pool”](#) on page 521.

See [“Edit or delete a volume pool”](#) on page 521.

Add a volume pool

Use this procedure to add a new volume pool.

To add a volume pool

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume pools** tab.
- 3 Select the **Add volume pool** button.
- 4 Specify the properties for the volume pool.
See [“Volume pool properties”](#) on page 522.
- 5 You can add volumes to the pool in the following ways:
 - Add new volumes to NetBackup.
 - Change the pool of existing volumes.

Edit or delete a volume pool

Edit a volume pool

Use this procedure to change the properties of a volume pool. The properties you can change include the pool type (scratch pool or catalog backup pool).

To edit a volume pool

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume pools** tab.
- 3 Select a pool in the **Volume pools** list.

- 4 Select **Edit**.
- 5 Change the attributes for the volume pool.
 See “[Volume pool properties](#)” on page 522.
- 6 Select **Update**.

Delete a volume pool

You cannot delete any of the following pools:

- A volume pool that contains volumes
- The **NetBackup** volume pool
- The **None** volume pool
- The default **CatalogBackup** volume pool
- The **DataStore** volume pool

To delete a volume pool

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume pools** tab.
- 3 Locate the volume pool that you want to delete. Ensure that the volume pool is empty. If the pool is not empty, change the pool name for any volumes in the pool.
- 4 Select the volume pool.
- 5 Select **Delete**.
- 6 Select **Confirm**.

Volume pool properties

You can specify various properties for a volume pool.

The following table describes the properties that you can configure for volume pools, either when you add a new pool or change an existing one.

Table 9-17 Volume pool properties

Property	Description
Catalog backup pool	Select this option to use this volume pool for catalog backups. This check box creates a dedicated catalog backup pool to be used for NBU-Catalog policies. A dedicated catalog volume pool facilitates quicker catalog restore times. Multiple catalog backup volume pools are allowed.
Description	a brief description of the volume pool.
Maximum number of partially full media	This property does not apply to the None pool, catalog backup pools, or scratch volume pools. Specifies the number of partially full media to allow in the volume pool for each of the unique combinations of the following in that pool: ■ Robot ■ Drive type ■ Retention level The default value is zero, which does not limit the number of full media that are allowed in the pool.
Prefer span to scratch	Specifies how NetBackup should select additional media when tape media operations span multiple media. When this parameter is set to <code>yes</code> (default) if a job spans to new media, NetBackup selects media from the scratch pool. NetBackup takes this action instead of using partially full media from the backup volume pool. When this parameter is set to <code>no</code>, NetBackup attempts to select partially full media from the backup volume pool to complete the specified operation. The <code>no</code> setting lets NetBackup use partially full media in the backup volume pool instead of always spanning to a scratch tape. Set the maximum number of partially full media option with the <code>vmppool -create</code> or the <code>vmppool -update</code> command.
Pool name	The Pool name is the name for the new volume pool. Volume pool names are case-sensitive and can be up to 20 characters.
Scratch pool	Specifies that the pool should be a scratch pool. It is recommended that you use a descriptive name for the pool and use the term <code>scratch pool</code> in the description. Add sufficient type and quantity of media to the scratch pool to service all scratch media requests that can occur. NetBackup requests scratch media when media in the existing volume pools are allocated for use. In NetBackup there can be only one scratch pool. This option is only available if no other pool is assigned as the scratch pool.

Managing volume groups

You can perform the following tasks to manage volume groups.

Moving a volume group

You can move a volume group from a robotic library to standalone storage or from standalone storage to a robotic library.

Moving a volume group changes only the residence information in the NetBackup database. You must move the volumes physically to their new locations.

To move a volume group

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, select the volume group that you want to move.
- 3 On the **Actions** menu, select **Move**.
- 4 In the **Move Volume Group** dialog box, specify the properties for the move.
You can only specify the properties that apply for the move type.

Property	Description
Destination	The destination for the move, as follows: ■ If you move the volume group from a robotic library, Standalone is selected as the destination. ■ If you move the volume group from standalone, Robot is selected as the destination.
Device host	The host that controls the robotic library.
Robot	The destination robotic library.
Volume group	The volume group to move. Displays "----" when you move standalone volumes.

- 5 After you move the volume group logically, physically move the volumes to their new locations.

Deleting a volume group

Use the following procedure to delete a volume group.

To delete a volume group

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the volumes list, verify that all of the volumes in the group are unassigned. You cannot delete the group until the application unassigns the volumes. If the **Time Assigned** column contains a value, the volume is assigned.
See [“About assigning and deassigning volumes”](#) on page 505.
- 3 Select a volume group in the right pane.
- 4 On the **Edit** menu, select **Delete**.
- 5 In the confirmation dialog box, confirm the action.
- 6 Remove the deleted volumes from the storage device.

About media sharing

Media sharing allows media servers to share media for write purposes (backups).

Media sharing provides the following benefits:

- Increases the utilization of media by reducing the number of partially full media.
- Reduces media-related expenses because fewer tape volumes are required and fewer tape volumes are vaulted (NetBackup Vault option).
- Reduces administrative overhead because you inject fewer scratch media into the robotic library.
- Increases the media life because tapes are mounted fewer times. Media are not repositioned and unmounted between write operations from different media servers.

Reducing media mounts requires appropriate hardware connectivity between the media servers that share media and the drives that can write to that media. Appropriate hardware connectivity may include Fibre Channel hubs or switches, SCSI multiplexors, or SCSI-to-fibre bridges.

You can configure the following media sharing:

- Unrestricted media sharing.
See [“Configure unrestricted media sharing”](#) on page 526.
- Media media sharing with server groups.
See [“Configure media sharing with a server group”](#) on page 526.\

Note: The access control feature of Sun StorageTek ACSLS controlled robots is not compatible with media sharing. Media sharing restricts volume access by the requesting hosts IP address. Use caution when you implement media sharing in an ACSLS environment.

Configure unrestricted media sharing

Unrestricted media sharing means that all NetBackup media servers and NDMP hosts in your NetBackup environment can share media for writing.

Note: Do not use both unrestricted media sharing and media sharing server groups. If you use both, NetBackup behavior is undefined.

To configure unrestricted media sharing

- 1 Open the NetBackup web UI.
- 2 On the left, select **Hosts > Host properties**.
- 3 Select the primary server.
- 4 If necessary, select **Connect**. Then select **Actions > Edit primary server**.
- 5 Select **Media**.
- 6 Select **Enable unrestricted media sharing for all media servers**.

If you allow unrestricted allow media sharing in your NetBackup environment, you do not need to create media sharing groups.

- 7 Select **Save**.

Configure media sharing with a server group

Media sharing with a server group restricts the sharing to members of the group.

See [“About NetBackup server groups”](#) on page 374.

[Table 9-18](#) outlines the process for configuring media sharing with a server group.

Note: Do not use both unrestricted media sharing and media sharing server groups. If you use both, NetBackup behavior is undefined.

Table 9-18 Configuring media sharing with a server group process overview

Step	Action	Description
Step 1	Ensure the appropriate connectivity between and among the media servers and robots and drives.	Beyond the scope of the NetBackup documentation.
Step 2	Configure the media sharing server group.	See “Add a server group” on page 374.
Step 3	Optionally, configure the volume pools for media sharing.	Set the Maximum number of partially full media property for those pools. See “Add a volume pool” on page 521. See “Edit or delete a volume pool” on page 521.
Step 4	Configure the backup policies that use the volume pools and media sharing groups.	Set the Policy volume pool and Media owner properties of the backup policies. See “Creating a backup policy” on page 657.

Inventorying robots

This chapter includes the following topics:

- [About robot inventory](#)
- [When to inventory a robot](#)
- [About showing a robot's contents](#)
- [Showing the media in a robot](#)
- [About comparing a robot's contents with the volume configuration](#)
- [Comparing media in a robot with the volume configuration](#)
- [About previewing volume configuration changes](#)
- [Previewing volume configuration changes for a robot](#)
- [About updating the NetBackup volume configuration](#)
- [Updating the NetBackup volume configuration with a robot's contents](#)
- [Robot inventory options](#)
- [About the vmphyinv physical inventory utility](#)

About robot inventory

Robot inventory is a logical operation that verifies the presence of media. (Robot inventory does not inventory the data on the media.)

After you physically add, remove, or move volumes in a robot, use a robot inventory to update the NetBackup volume configuration.

The following table describes the robot inventory options for the robotic libraries that contain barcode readers and contain barcoded media.

Table 10-1 Robot inventory options

Inventory option	Description
Show contents	Queries the robot for its contents and displays the media in the selected robotic library; does not check or change the NetBackup database. See “About showing a robot's contents” on page 531. For the robotic libraries without barcode readers (or that contain media without barcodes), you can only show the contents of a robot. However, more detailed information is required to perform automated media management. Use the <code>vmphyinv</code> physical inventory utility to inventory such robots. See “About the vmphyinv physical inventory utility” on page 540.
Compare contents with volume configuration	Queries the robot for its contents and compares the contents with the contents of the NetBackup database. Does not change the database. See “About comparing a robot's contents with the volume configuration” on page 534.
Preview volume configuration changes	Queries the robot for its contents and compares the contents with the contents of the NetBackup database. If differences exist, it is recommended to change to the NetBackup volume configuration. See “About previewing volume configuration changes” on page 535.
Update volume configuration	Queries the robot for its contents; if necessary, updates the database to match the contents of the robot. If the robot contents are the same as the NetBackup database, no changes occur. See “About updating the NetBackup volume configuration” on page 537.

When to inventory a robot

The following table describes the criteria to use to determine when to inventory a robot and which options to use for the inventory.

Table 10-2 Robot inventory criteria

Action	Inventory option to use
To determine the contents of a robot	Use the Show contents option to determine the media in a robot and possibly their barcode numbers. See “Showing the media in a robot” on page 533.

Table 10-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
To determine if volumes were moved physically within a robot	For the robots with barcode readers and the robots that contain media with barcodes, use the Compare contents with volume configuration option. See “Comparing media in a robot with the volume configuration” on page 535.
To add new volumes to a robot (a new volume is one that does not have a NetBackup media ID)	For any robot NetBackup supports, use the Update volume configuration option. The update creates media IDs (based on barcodes or a prefix that you specify). See “Updating the NetBackup volume configuration with a robot's contents” on page 538.
To determine whether new media have barcodes before you add them to NetBackup	Use the Preview volume configuration changes option, which compares the contents of the robot with the NetBackup volume configuration information. After you examine the results, use the Update volume configuration option to update the volume configuration if necessary. See “Updating the NetBackup volume configuration with a robot's contents” on page 538.
To insert existing volumes into a robot (an existing volume is one that already has a NetBackup media ID)	If the robot supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic location. NetBackup also updates the robot host, robot type, robot number, and slot location. Specify the volume group to which the volume is assigned. See “Updating the NetBackup volume configuration with a robot's contents” on page 538. If the robot does not support barcodes or the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility. See “About moving volumes” on page 517. See “About the vmphyinv physical inventory utility” on page 540.
To move existing volumes between robotic and standalone (an existing volume is one that already has a NetBackup media ID)	If the robotic library supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic or standalone location. See “Updating the NetBackup volume configuration with a robot's contents” on page 538.

Table 10-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
To move existing volumes within a robot (an existing volume is one that already has a NetBackup media ID)	If the robot supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new slot location. See “Updating the NetBackup volume configuration with a robot's contents” on page 538. If the robot does not support barcodes or if the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility. See “About moving volumes” on page 517. See “About the vmphyinv physical inventory utility” on page 540.
To move existing volumes from one robot to another (an existing volume is one that already has a NetBackup media ID)	If the robotic library supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the NetBackup volume configuration information. See “Updating the NetBackup volume configuration with a robot's contents” on page 538. If the robots do not support barcodes or the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility. See “About moving volumes” on page 517. See “About the vmphyinv physical inventory utility” on page 540. For either operation, perform the following updates: ■ First move the volumes to standalone ■ Then move the volumes to the new robot If you do not perform both updates, NetBackup cannot update the entries and writes an "Update failed" error.
To remove existing volumes from a robot (an existing volume is one that already has a NetBackup media ID)	For any robot NetBackup supports, use the Update volume configuration option to update the NetBackup volume configuration information. See “Updating the NetBackup volume configuration with a robot's contents” on page 538.

About showing a robot's contents

Show contents inventories the selected robotic library and generates a report. This operation does not check or change the NetBackup database. Use this option to determine the contents of a robot.

The contents that appear depend on the robot type.

The following table describes the report contents.


Note: On UNIX: If a volume is mounted in a drive, the inventory report lists the slot from which the volume was moved to the drive.

Table 10-3 Show contents description

Robot and media	Report contents
The robot has a barcode reader and the robot contains media with barcodes.	Shows if each slot has media and lists the barcode for the media.
The robot does not have a barcode reader or the robot contains media without barcodes.	Shows if each slot has media.
API robot.	Shows a list of the volumes in the robot. See “About inventory results for API robots” on page 532.

The following figure shows an example of the report.

Figure 10-1 Show contents report



See [“Showing the media in a robot”](#) on page 533.

About inventory results for API robots


The following table describes the contents of the robot inventory for the API robots.

Table 10-4 API robot report contents

Robot type	Report contents
ACS	The results, received from ACS library software, show the following: ■ The ACS library software volume ID. The NetBackup media ID corresponds to the ACS library software volume ID. ■ The ACS media type. ■ The NetBackup Media Manager media type. ■ The mapping between the ACS library software media type and the corresponding NetBackup Media Manager media type (without considering optional barcode rules).

The following figure shows the results for an ACS robot; the results for other API robots are similar.

Figure 10-2 Show contents report (API robot)



Media ID	ACS Media Type	MM Media Type
000026	STK1R	HCART
000042	STK1R	HCART
000043	STK1R	HCART
000044	STK1R	HCART
000045	STK1R	HCART
000046	STK1R	HCART
000047	STK1R	HCART
000049	STK1R	HCART
000051	STK1R	HCART

Showing the media in a robot

Use the following procedure to show the media that is in a robot.

See [“About robot inventory”](#) on page 528.

See [“Robot inventory options”](#) on page 539.

To show the media in a robot

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Robots** tab.
- 3 Select the robot you want to inventory.

- 4 Select **Inventory robot**.
- 5 Go to the **Inventory operation** settings and select **Show contents**.
- 6 Select **Start** to begin the inventory.

About comparing a robot's contents with the volume configuration

Compare contents with volume configuration compares the contents of a robotic library with the contents of the NetBackup database. Regardless of the result, the database is not changed.

Table 10-5 Compare contents description

Robot and media	Report contents
The robot can read barcodes	The report shows the differences between the robot and the NetBackup database
The robot cannot read barcodes	The report shows only whether a slot contains a volume If the media cave barcodes, this operation is useful for determining if volumes have been physically moved within a robot.
For API robots	The media ID and media type in the NetBackup database are compared to the information that is received from the vendor's robotic library software.

If the results show that the NetBackup database does not match the contents of the robotic library, perform the following actions:

- Physically move the volume.
- Update the NetBackup database. Use **Move** or use the **Update volume configuration** option.

See [“About updating the NetBackup volume configuration”](#) on page 537.

The following figure shows a sample compare report.

Figure 10-3 Compare contents report (API robot)

Results						
Robot Contents			Volume Configuration			
Media ID	ACS Media Type	MM Media Type	Media ID	Media Type	Mismatch?	
=====			=====		=====	
000026	STKLR	HCART	000026	HCART		
000042	STKLR	HCART	000042	HCART		
000043	STKLR	HCART	000043	HCART		
000044	STKLR	HCART	000044	HCART		
000045	STKLR	HCART	000045	HCART		

Clear Results

See [“Comparing media in a robot with the volume configuration”](#) on page 535.

Comparing media in a robot with the volume configuration

Use the following procedure to compare the media in a robot with the NetBackup database.

- See [“About robot inventory”](#) on page 528.
- See [“Robot inventory options”](#) on page 539.

To compare media in a robot with the volume configuration

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Robots** tab.
- 3 On the **Actions** menu, select **Inventory robot**.
- 4 Select the **Device host** and the **Robot** that you want to inventory.
- 5 Select **Compare contents with volume configuration**.
- 6 Select **Start** to begin the inventory.

About previewing volume configuration changes

Use this option to preview the changes before you update the NetBackup database. This option lets ensure that all new media have barcodes before you add them to the NetBackup database.

Note: If you preview the configuration changes first and then update the NetBackup database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the NetBackup database, to the barcode rules, and so on.

```

-----
September 26, 2024 15:31:45 GMT
Robot: TLD(1) on server.domain.com
Operation: Preview of inventory and update
Total 1 robot inventory volume config results
-----
Generating list of recommended changes ...

Proposed Change(s) to Update the Volume Configuration
=====
Logically move media ID 1A007B (barcode 001A007B) from standalone to slot 264

See “Updating the NetBackup volume configuration with a robot's contents”
on page 538.
```

Previewing volume configuration changes for a robot

Use the procedure in this topic to preview any volume configuration changes for a robot.

See [“About previewing volume configuration changes”](#) on page 535.

See [“Robot inventory options”](#) on page 539.

To preview the volume configuration changes for a robot

- 1 If necessary, insert new volumes into the robotic library.
- 2 Open the NetBackup web UI.
- 3 On the left, select **Storage > Tape storage**. Then select the **Robots** tab.
- 4 Select the robot that you want to inventory.
- 5 Select **Inventory robot**.
- 6 Select the device host and the robot that you want to inventory.

7 Select **Preview volume configuration changes**.

Note: If you preview the configuration changes first and then update the NetBackup database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the NetBackup database, to the barcode rules, and so on.

- 8 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, adjust the settings in **Advanced options**.
See [“About configuring media name and attribute rules”](#) on page 473.
- 9 To inject any media that is in the media access port before the preview operation, select **Empty media access port prior to update**.
- 10 Select **Start** to begin the inventory preview.

About updating the NetBackup volume configuration

The **Update volume configuration** robot inventory option updates the database to match the contents of the robot. If the robot contents are the same as the NetBackup database, no changes occur.

For a new volume (one that does not have a NetBackup media ID), the update creates a media ID. The media ID depends on the rules that are specified.

See [“Robot inventory options”](#) on page 539.

For API robots, the update returns an error if the volume serial number or the media ID contain unsupported characters.


For robots without barcode readers, the new media IDs are based on a media ID prefix that you specify. Similarly, for volumes without readable barcodes, the new media IDs are based on a media ID prefix that you specify.

[Figure 10-4](#) is an example for an ACS robot. Results for other API robots are similar.

Robot inventory update returns an error if it encounters unsupported characters in the volume serial number or media identifier from API robots.

See [“Volume update prerequisites”](#) on page 538.

Figure 10-4 Update volume configuration for API robot report



See [“Updating the NetBackup volume configuration with a robot's contents”](#) on page 538.

Volume update prerequisites

The following are the robot prerequisites and media prerequisites for updating the volume configuration:

- The robotic library must read barcodes.
- Volumes in the library must have readable barcodes.

You can check the barcode capabilities of the robotic library and the volumes by comparing the robot contents with the NetBackup volume configuration.

See [“Comparing media in a robot with the volume configuration”](#) on page 535.

If the robotic library does not support barcodes or the volumes do not have readable barcodes, save the results of the compare operation. The results can help you determine a media ID prefix if you use the **Media Settings** tab of the **Advanced Options** dialog box to assign a prefix.

Updating the NetBackup volume configuration with a robot's contents

Use the procedure in this topic to update the NetBackup database with the contents of a robot.

See [“About updating the NetBackup volume configuration”](#) on page 537.

See [“Robot inventory options”](#) on page 539.

To update the volume configuration with a robot's contents

- 1 If necessary, insert new volumes into the robotic library.
- 2 Open the NetBackup web UI.
- 3 On the left, select **Storage > Tape storage**. Then select the **Robots** tab.
- 4 Select **Actions > Inventory robot**.
- 5 Select the device host and the robot that you want to inventory.
- 6 Select **Update volume configuration**.

Note: If you preview the configuration changes first and then update the NetBackup database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the NetBackup database, to the barcode rules, and so on.

- 7 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, select **Advanced options**.
See [“About configuring media name and attribute rules”](#) on page 473.
- 8 To inject any media that is in the media access port before the update operation, select **Empty media access port prior to update**.
- 9 Select **Start** to begin the inventory update.

Robot inventory options

The following table shows the robot inventory options.

Table 10-6 Robot inventory options

Option	Description
Advanced options	The Advanced options menu is active if Preview volume configuration changes or Update volume configuration is selected. This menu includes additional configuration options. See “About configuring media name and attribute rules” on page 473.
Device host	The Device host option is the host that controls the robot.

Table 10-6 Robot inventory options (*continued*)

Option	Description
Empty media access port prior to update	The Empty media access port prior to update operation is active only for the robots that support that function. To inject volumes in the robot's media access port into the robot before you begin the update, select Empty media access port prior to update. The volumes to be injected must be in the media access port before the operation begins. If you select Empty media access port prior to update and the media access port is empty, you are not prompted to place volumes in the media access port. Note: If you use NetBackup to eject volumes from the robot, remove the volumes from the media access port before you begin an inject operation. Otherwise, if the inject port and eject port are the same, the ejected volumes may be injected back into the robotic library.
Robot	Use the Robot option to select a robot to inventory.
Show contents	Displays the media in the selected robotic library; does not check or change the NetBackup database. See “About showing a robot's contents” on page 531.
Compare contents with volume configuration	Compares the contents of a robotic library with the contents of the NetBackup database but does not change the database. See “About comparing a robot's contents with the volume configuration” on page 534.
Preview volume configuration changes	Compares the contents of a robotic library with the contents of the NetBackup database. If differences exist, it is recommended to change to the NetBackup volume configuration. See “About previewing volume configuration changes” on page 535.
Update volume configuration	Updates the database to match the contents of the robot. If the robot contents are the same as the NetBackup database, no changes occur. See “About updating the NetBackup volume configuration” on page 537.

About the vmphyinv physical inventory utility

For the following robotic libraries, NetBackup reports only the presence of media:

- For the robots without barcode readers
- For the robots that contain media without barcodes

More detailed information is required to perform automated media management. For such robots, use the `vmphyinv` physical inventory utility.

The `vmphyinv` physical inventory utility inventories non-barcoded tape libraries by performing the following actions:

- Mounts each tape
 - Reads the tape header
 - Identifies the tape in each slot
 - Updates the NetBackup volume configuration
- Use the `vmphyinv -verbose` option to display more information about the suggested changes. The `-verbose` option shows the number of drives available, the contents of each tape, if the media is a catalog tape. (The media format column of the summary contains NetBackup database for NetBackup catalog tapes.)
- This verbose information is written to `stderr`. To save the information, redirect `stderr` to a file.

`vmphyinv` is a command-line utility. Additional information about the syntax of the `vmphyinv` command is available.

For more information, see the [NetBackup Commands Reference Guide](#).

Table 10-7 `vmphyinv` features, requirements, restrictions, and when to use

Item	Description
vmphyinv features	The <code>vmphyinv</code> utility has the following features: ■ Can be run from any primary server, media server, or SAN media server. ■ Can be used with barcoded tape libraries because it verifies the contents of the media. ■ Recognizes the NetBackup tape formats. ■ Supports the remote administration. You do not need to run <code>vmphyinv</code> from the host to which the drives are attached. ■ Tries to use multiple drives in a robot even if the drives are attached to different hosts. ■ Works with shared drives (NetBackup Shared Storage Option). ■ Supports all supported SCSI-based robot types. ■ Can be used to inventory a single media in a standalone drive. Use the <code>-u</code> option or the <code>-n</code> option to specify the drive; the drive must contain media and it must be ready.

Table 10-7 `vmphyinv` features, requirements, restrictions, and when to use
(continued)

Item	Description
vmphyinv requirements and restrictions	The <code>vmphyinv</code> utility has the following requirements and restrictions: ■ It cannot distinguish between the volume records based on the application type. ■ When you move the media from robotic drives to standalone drives, you cannot specify a new volume group for the media.
When to use <code>vmphyinv</code>	Use <code>vmphyinv</code> to update the NetBackup database for NetBackup in the following cases: ■ You want to inventory a robot that does not have a barcode reader or that contains non-barcoded media. ■ You insert new media into a robotic library and no NetBackup volume records correspond to the media. Use the slot range or list option of <code>vmphyinv</code> to perform the inventory operation. You do not need to add volume records to the NetBackup database. ■ You insert some media that have unknown media IDs or globally unique identifiers (GUIDs) into a robot. For example, you insert 10 media from a different tape library in slots 11 to 20. You do not know the IDs on the tapes. Use the slot range or list option of <code>vmphyinv</code> to perform the inventory operation. The <code>vmphyinv</code> utility mounts the media, reads the tape header, determines the media ID, and adds media records to the NetBackup database. ■ Some of the media are misplaced and the NetBackup database does not reflect the correct physical location of these media. Inventory the robot or inventory a subset of media in the robot by using options in <code>vmphyinv</code>.

How `vmphyinv` performs a physical inventory

For a physical inventory, the `vmphyinv` utility performs the following sequence of operations:

- Obtains a list of drives to mount the media
See [“About the `vmphyinv` list of drives”](#) on page 543.
- Obtains a list of media to mount
See [“About the media that `vmphyinv` mounts”](#) on page 543.
- Mounts the media and reads the tape headers
See [“How `vmphyinv` mounts the media and reads the tape header”](#) on page 544.
- Updates the NetBackup database
See [“How `vmphyinv` updates the NetBackup database”](#) on page 545.

About the vmphyinv list of drives

The list of drives the `vmphyinv` utility uses to mount the media is obtained from the NetBackup database. The drives do not need to be configured locally.

You cannot specify which drives to use. However, you can specify the maximum number of drives to use, which lets you reserve drives for NetBackup backup or restore operations. Specify the number of drives by using the `-drv_cnt drive_count` option.

About the media that vmphyinv mounts

The `vmphyinv` command accepts several options for the media to be mounted, as follows:

- NetBackup robot number (`-rn robot_number`).
 The `vmphyinv` utility obtains a list of volume records for that robot and inventories each of the media in the list.
 To use this option, the NetBackup configuration must contain a volume record that corresponds to the robot number in the NetBackup database for the robot.

- NetBackup robot number with filter options.
 If you do not want to inventory all of the media in a robot, specify a subset of the media by using filter options. Some filter options are volume pool, volume group, or slot range. To use these options, NetBackup volume records must exist.

The following are some filter examples.

<code>vmphyinv -rn 4 -pn bear</code>	Mounts the media only in robot 4 and in the volume pool bear.
<code>vmphyinv -rn 2 -v moon</code>	Mounts the media in robot 2 and in the volume group moon.
<code>vmphyinv -rn 1 -rc1 2 -number 3</code>	Mounts the media in robot 1 and slot range 2 to 4.
<code>vmphyinv -rn 5 -pn NetBackup -v mars -rc1 2 -number 6</code>	Mounts the media in robot 5, slot range 2 to 7, in volume group mars, and in the NetBackup volume pool.

- NetBackup robot number and a list of media that belong to a specific robot.
 For example, if the `-rn robot_number` and `-ml A00001:A00002:A00003` options are specified, only the three specified media are inventoried. If any of these media do not belong to the specified robot, the media are skipped and are not inventoried. To use this option, NetBackup volume records must exist.

- NetBackup robot number and a slot range or list.
Sometimes, media from a different robot or some other source are moved to a robot and the media ID on the tape is unknown. In these cases, specify a slot range option or list option.
With these options, the NetBackup volume record does not need to exist in the NetBackup database. However, you must specify the density (using the `-d` option).

Note: For a robot that supports multiple media types, specify the density carefully. If you specify the incorrect density, `vmphyinv` cannot complete the mount and permanent drive failure can occur.

The following are some filter examples.

<code>vmphyinv -rn 1 -slot_range 2 10 -d dlt</code>	Mounts the media in slot range 2 to 10 in robot 1.
<code>vmphyinv -rn 0 -slot_list 3:4:5 -d dlt</code>	Mounts the media in slots 3, 4, and 5 in robot 0.
<code>vmphyinv -rn 2 -slot_range 2 4 -slot_list 5:6:7 -d dlt</code>	Mounts the media in slots 2, 3, 4, 5, 6, and 7 in robot 2.

See [“About the `vmphyinv` physical inventory utility”](#) on page 540.

How `vmphyinv` mounts the media and reads the tape header

The following sequence of operations explains the mount process:

- The `vmphyinv` utility contacts the NetBackup Volume Manager, `vmid`, on the local host or remote host depending on where the drive is attached.
- The NetBackup Volume Manager starts a process, `opr`.
- The `vmphyinv` utility communicates with `opr` and sends the mount request to `opr`. After `opr` receives the request, it issues a mount request to `ltid`.
- The `vmphyinv` utility reads the tape header to determine the recorded media ID or globally unique identifier (GUID).

Note: The default mount timeout is 15 minutes. Specify a different mount time by using the `-mount_timeout` option.

See [“About the media that `vmphyinv` does not recognize”](#) on page 545.

See [“How `vmphyinv` processes cleaning media”](#) on page 545.

About the media that `vmphyinv` does not recognize

If the media is not NetBackup media, the media is unmounted and the next media is mounted. `vmphyinv` does not generate a new record in the NetBackup database. To generate volume records for that media, use the `vmupdate` command.

How `vmphyinv` processes cleaning media

If the following conditions are all true, `vmphyinv` does not try to mount the media and the next media in the list is mounted:

- You do not specify the `vmphyinv` slot range or list option.
- The robot contains cleaning media.
- The media type is specified as cleaning media in the volume record (such as `hcart2_clean` or `dlt_clean`).

If the robot contains cleaning media and any of the following conditions are true, `vmphyinv` tries to determine if the media is cleaning media:

- You use the slot range or list option and the media type of volume record in the NetBackup database is not a cleaning media type.
- You use the slot range or list option, and the NetBackup database does not contain a volume record that corresponds to the cleaning media.
- You do not use the slot range or list option, and the NetBackup database does not contain a volume record that corresponds to the cleaning media.

The `vmphyinv` utility tries to determine if the media is cleaning media. It uses the SCSI parameters (sense keys, tape alert flags, and physical (SCSI) media types) returned by the robot. If `vmphyinv` cannot determine if the media is cleaning media, it tries to mount the media until the mount request times out.

Note: NetBackup may not detect the presence of cleaning media for all drives. Some drives report the presence of cleaning media in a manner NetBackup cannot read.

How `vmphyinv` updates the NetBackup database

After all of the media are mounted and the tape headers are read, `vmphyinv` displays a list of recommended changes. Accept or reject the changes. If you accept the changes, `vmphyinv` updates the NetBackup database.

Table 10-8 vmphyinv criteria and actions

Criteria or action	Description
The <code>vmphyinv</code> update criteria	For valid media types, <code>vmphyinv</code> performs the following actions: ■ Changes the residence fields and description fields of any NetBackup media record if those fields do not match the media header. ■ Conditionally changes the media type of an unassigned NetBackup volume record. The media type is changed only if the new media type belongs to the same family of media types as the old media type. For example, the media type DLT can only be changed to DLT2 or DLT3. ■ Never changes the volume pool, media type, and ADAMM_GUID of an assigned record. ■ Never unassigns an assigned NetBackup volume.
How <code>vmphyinv</code> updates NetBackup media	The <code>vmphyinv</code> utility searches the NetBackup database. It checks if the media ID from the tape is present in the media ID field of any record in the NetBackup database. If the media ID exists, <code>vmphyinv</code> updates the NetBackup volume record that corresponds to the media ID. If the media ID does not exist, <code>vmphyinv</code> creates a new NetBackup volume record that corresponds to the NetBackup media.
<code>vmphyinv</code> error cases	The <code>vmphyinv</code> utility may not be able to update the NetBackup database correctly in the following cases. These cases are reported as errors. If any of the following cases are encountered, you must intervene to continue: ■ Duplicate media IDs are found. Two or more media in the same robot have the same media ID. ■ A NetBackup volume record that belongs to a different robot is found. It contains the same media ID as the media ID read from the tape header. ■ The media type, media GUID, or volume pool of an assigned volume record needs to be changed. ■ The barcode of an existing volume record needs to be changed.

Configuring storage units

This chapter includes the following topics:

- [About storage](#)
- [Creating a storage unit](#)
- [About storage unit settings](#)
- [Overview of universal shares](#)

About storage

The data that is generated from a NetBackup job is recorded into a type of storage that NetBackup recognizes.

NetBackup recognizes the following storage configurations, all of which are configured in **Storage**:

Storage units

A storage unit is a label that NetBackup associates with physical storage. The label can identify a robot, a path to a volume, or a disk pool. Storage units can be included as part of a storage unit group or a storage lifecycle policy.

See [“Creating a storage unit”](#) on page 548.

Storage unit groups

Storage unit groups let you identify multiple storage units as belonging to a single group. The NetBackup administrator configures how the storage units are selected within the group when a backup or a snapshot job runs.

See [“About storage unit groups”](#) on page 590.

Storage lifecycle policies

Storage lifecycle policies let the administrator create a storage plan for all of the data in a backup or snapshot.

See [“About storage lifecycle policies”](#) on page 601.

Creating a storage unit

A storage unit is a label that NetBackup associates with physical storage. The label can identify a robot, a path to a volume, or a disk pool.

Storage unit creation is part of several other wizards. However, a storage unit can be created directly from the **Storage** utility.

[Figure 11-1](#) shows the different storage unit types and the option that needs to be installed, if necessary.

Figure 11-1 Storage unit types

Media Manager	NDMP	Disk
Points to a robot or a standalone drive	Points to an NDMP host (NDMP Option)	AdvancedDisk Points to a disk pool (storage directly attached to a media server). (Data Protection Optimization Option) BasicDisk Points to a directory. Cloud Storage (VendorName) Points to a disk pool of the type identified by VendorName where VendorName can be the name of a cloud storage provider. OpenStorage (StorageName) Points to a disk pool of the type identified by StorageName. (OpenStorage Disk Option) PureDisk Points to a disk pool. (Data Protection Optimization Option)

To create a storage unit

- 1 In the **NetBackup Administration Console**, select the **Storage** utility.
 Some storage unit types can also be created in the NetBackup web UI. On the left, click **Storage > Storage units**.
- 2 Select **Actions > New > New Storage Unit**.
- 3 Enter a **Storage unit name**.
 See [“NetBackup naming conventions”](#) on page 1054.
- 4 Select the **Storage unit type**. The selection specifies the type of storage that the storage unit uses: **Media Manager**, **Disk**, or **NDMP**.
- 5 For disk storage units:
 - Select a disk type from the **Disk type** drop-down menu.
 The **Disk type** identifies the type of storage unit destination:

AdvancedDisk storage unit	The destination is a disk pool.
BasicDisk storage unit	The destination is a path to a volume on a host.
Cloud storage unit	The destination is a disk pool of the type that includes a VendorName string. VendorName can be the name of a cloud storage provider. The values also can contain a _crypt suffix (for example, Amazon_crypt). The _crypt suffix indicates encrypted storage.
NDMP storage	The destination is an NDMP host. The NDMP protocol is used to perform backups and recoveries.
OpenStorage storage unit	The destination is a disk pool of the type that includes a VendorName string. The vendor supplies the StorageName string. For DataDomain, you can use the WORM property of the disk. WORM is the acronym for Write Once Read Many. If the Use WORM option is set, data can be written to the associated media only once, but it can be read multiple times.
PureDisk storage unit	The destination is a Media Server Deduplication Pool. See the NetBackup Deduplication Guide.
 - Select the disk pool for this storage unit.
 See [“Select disk pool storage unit setting”](#) on page 572.

- Select a media server in the **Media server** drop-down menu. The selection indicates that the media server has permission to write to the storage unit.
 - **Absolute pathname to directory** or **Absolute pathname to volume setting**.
See [“Absolute pathname to directory or absolute pathname to volume setting for storage units”](#) on page 559.
 - **Maximum concurrent jobs**
See [“Maximum concurrent jobs storage unit setting”](#) on page 563.
 - **Reduce fragment size**
See [“Reduce fragment size storage unit setting”](#) on page 571.
 - **High water mark**
See [“High water mark storage unit setting”](#) on page 561.
 - **Low water mark**
See [“Low water mark storage unit setting”](#) on page 562.
 - **Enable block sharing**
See [“Enable block sharing storage unit setting”](#) on page 561.
 - **Enable Temporary staging area**
See [“Enable temporary staging area storage unit setting”](#) on page 573.
- 6 For Media Manager storage units, data is written to tape robots and standalone tape drives:
- Select a storage device from the **Storage Device** drop-down menu.
 - Select a media server in the **Media server** drop-down menu. The selection indicates that the media server has permission to write to the storage unit.
 - **Maximum concurrent write drives**
See [“Maximum concurrent write drives storage unit setting”](#) on page 563.
 - **Enable multiplexing**
See [“Enable multiplexing storage unit setting”](#) on page 561.
 - **Reduce fragment size**
See [“Reduce fragment size storage unit setting”](#) on page 571.
- 7 Click **OK** to save the storage unit configuration.

Creating a storage unit by copying a storage unit

You can copy a storage unit to create a new storage unit with the same settings. This option is not available for **OST** storage type.

See [the section called “Copy a disk storage unit”](#) on page 551.

See [the section called “Copy a tape storage unit”](#) on page 551.

Copy a disk storage unit

To copy a disk storage unit

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Storage units**. Select the **Storage units** tab.
- 3 Select the storage unit that you want to copy and select the **Copy storage unit** button.
- 4 Type a unique name for the new storage unit. For example, describe the type of storage. Use this name to specify a storage unit for policies and schedules.
See [“NetBackup naming conventions”](#) on page 1054.
- 5 Edit the other properties and disk pool as necessary.
See [“About storage unit settings”](#) on page 559.
- 6 After reviewing the changes, select the **Save** button.

Copy a tape storage unit

To copy a tape storage unit

- 1 On the left, click on the **Storage > Storage units**.
- 2 Select the tape storage unit that you want to copy and select the **Copy storage unit** button.

Note that the storage unit name is appended with “_Copy”.
- 3 Make any changes that you want and select the **Save** button.

Editing storage unit settings

Only make changes to a storage unit during periods when no backup activity is expected. This way backups are not affected for the policies or protection plans that use the affected storage units.

To edit storage unit settings

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Storage units**. Select the **Storage units** tab.
- 3 Select the storage unit that you want to edit.
- 4 Click **Edit** and make the required changes.

See [“About storage unit settings”](#) on page 559.

For example, you can edit the following settings:

- The basic properties of the storage unit.
- Additional properties
- Media servers
- Staging schedule

To edit a tape storage unit

- 1 On the left, click **Storage > Storage units**.
- 2 On the tape storage unit list, click on the tape storage unit you want to edit.
- 3 Click **Edit** and make the required changes. Click **Save** after making the changes.

Deleting storage units

To delete a storage unit from a NetBackup configuration means to delete the label that NetBackup associates with the physical storage.

Deleting a storage unit does not prevent files from being restored that were written to that storage unit. (As long as the storage was not physically removed and the backup image has not expired.)

To delete a storage unit

- 1 Open the NetBackup web UI.
- 2 Use the **Catalog** utility to expire any images that exist on the storage unit. This action removes the image from the NetBackup catalog.

See [“Expiring backup images”](#) on page 933.

- Do not manually remove images from a storage unit.
- After the images are expired, they cannot be restored unless the images are imported.

See [“About importing backup images”](#) on page 935.

NetBackup automatically deletes any image fragments from a disk storage unit or a disk pool. This deletion generally occurs within seconds of expiring an image. However, to make sure that all of the fragments are deleted, confirm that the directory is empty on the storage unit.

- 3 On the left, select **Storage > Storage units**. Select the **Storage units** tab..
- 4 Select the storage unit that you want to delete.

- 5 Select **Delete > Yes**.
- 6 Modify any policy that uses a deleted storage unit to use another storage unit.
If a storage unit points to a disk pool, you can delete the storage unit without affecting the disk pool.

Tape storage unit considerations

To create a storage unit of a tape robot or a standalone tape drive, when you add a storage unit select **Drives and robots** as the **Storage type**.

See [“About storage unit settings”](#) on page 559.

When NetBackup sends a job to a **Tape** storage unit, a request is made to mount the volume in a drive.

If a standalone drive does not contain media or if a required volume is not available to a robot, a mount request appears in the **Pending requests**. (In the NetBackup web UI, open **Tape storage > Device monitor**). An operator can then find the volume, mount it manually, and assign it to the drive.

Take the following items into consideration when you add a storage unit for drives and robots:

- Where to add the storage unit depends on which version of NetBackup is in use.
 - Add the storage unit to the primary server. Specify the media server where the drives attach.
 - If using NetBackup Server, add the storage unit to the primary server where the drives attach. The robotic control must also attach to that server.
- The number of storage units that you must create for a robot depends on the robot's drive configuration.
 - Drives with identical densities must share the same storage unit on the same media server. If a robot contains two drives of the same density on the same media server, add only a single storage unit for the robot. Set the **Maximum concurrent write drives** setting to 2.
See [“Maximum concurrent write drives storage unit setting”](#) on page 563.
 - Drives with different densities must be in separate storage units. Consider an STK SL500 library that is configured as a Tape Library DLT (TLD). It can have both half-inch cartridge and DLT drives. Here, you must define a separate storage unit for each density.
 - If a robot's drives and robotic control attach to different NetBackup servers, specify the server where the drives attach as the media server. Always specify the same robot number for the drives as is used for the robotic control.

Disk storage unit considerations

NetBackup permits the creation of an unlimited number of disk storage units.

[Table 11-1](#) describes the different disk types that NetBackup can use as disk media.

Table 11-1 Disk media descriptions

Type of disk storage unit	Description
Media Server Deduplication Pool	A Media Server Deduplication Pool disk type storage unit is used for deduplicated data for a Media Server Deduplication Pool (MSDP), MSDP Cloud, or an MSDP volume group (MVG).
AdvancedDisk	An AdvancedDisk disk type storage unit is used for a dedicated disk that is directly attached to a NetBackup media server. An AdvancedDisk selection is available only when the Data Protection Optimization Option is licensed. NetBackup assumes the exclusive ownership of the disk resources that comprise an AdvancedDisk disk pool. If the resources are shared with other users, NetBackup cannot manage disk pool capacity or storage lifecycle policies correctly. For AdvancedDisk, the NetBackup media servers function as both data movers and storage servers. See the NetBackup AdvancedDisk Storage Solutions Guide.
Cloud Storage	A Cloud Storage disk type storage unit is used for storage in a cloud, usually provided by a third-party vendor. The actual name of the disk type depends on the cloud storage vendor. A Cloud Storage selection is available only when the Data Protection Optimization Option is licensed. The cloud storage provided by storage vendor partners is integrated into NetBackup via the API. A vendor host on the Internet is the storage server. The NetBackup media servers function as the data movers. See the NetBackup Cloud Administrator's Guide.
OpenStorage	An OpenStorage disk type storage unit is used for disk storage, usually provided by a third-party vendor. The actual name of the disk type depends on the vendor. An OpenStorage selection is available only when the OpenStorage Disk Option is licensed. The storage provided by storage vendor partners is integrated into NetBackup via the API. The storage host is the storage server. The NetBackup media servers function as the data movers. The storage vendor's plug-in must be installed on each media server that functions as a data mover. The logon credentials to the storage server must be configured on each media server. See the NetBackup OpenStorage Solutions Guide for Disk.

Table 11-1 Disk media descriptions (*continued*)

Type of disk storage unit	Description
Cloud Connector	An Cloud Connector disk type storage unit is used for backups to cloud object storage, with no deduplication. See the NetBackup for Cloud Object Store Administrator's Guide.
BasicDisk	A BasicDisk type storage unit consists of a directory on a locally-attached disk or a network-attached disk that is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory. Notes about the BasicDisk type storage unit: - Do not include the same volume or file system in multiple BasicDisk storage units. - BasicDisk storage units cannot be used in a storage lifecycle policy.

Not all settings are available on each disk storage unit type.

See [“About storage unit settings”](#) on page 559.

Note: It is recommended that you do not impose quotas on any file systems that NetBackup uses for disk storage units. Some NetBackup features may not work properly when file systems have quotas in place. (For example, the capacity-managed retention selection in storage lifecycle policies and staging to storage units.)

About the disk storage model

The NetBackup model for disk storage accommodates all Enterprise Disk Options. That is, it is the model for all disk types except for the BasicDisk type.

The following items describe components of the disk storage model:

Data mover

An entity that moves data between the primary storage (the NetBackup client) and the storage server. NetBackup media servers function as data movers.

Depending on the disk option, a NetBackup media server also may function as a storage server.

Storage server

An entity that writes data to and reads data from the disk storage. A storage server is the entity that has a mount on the file system on the storage.

Depending on the NetBackup option, the storage server is one of the following:

- A computer that hosts the storage. The computer may be embedded in the storage device.
- A storage vendor's host on the Internet that exposes cloud storage to NetBackup. Alternatively, private cloud storage can be hosted within your private network.
- A NetBackup media server that hosts storage.

Disk pool

A collection of disk volumes that are administered as an entity. NetBackup aggregates the disk volumes into pools of storage (a disk pool) you can use for backups.

A disk pool is a storage type in NetBackup. When you create a storage unit, you select the disk type and then you select a specific disk pool.

Configure the NetBackup service credentials for CIFS storage and disk storage units

For Common Internet File System (CIFS) storage with AdvancedDisk and BasicDisk storage units, the NetBackup Client Service and the NetBackup Remote Manager and Monitor Service services on Windows computers must use the same account credentials. If the account credentials are not configured properly, NetBackup marks all CIFS AdvancedDisk and BasicDisk storage units that use the UNC naming convention as DOWN.

To configure service credentials for CIFS storage and disk storage units

- 1 On the media servers that have a file system mount on the CIFS storage, configure the account and the credentials.

The account must be the same account that the Windows operating system uses for read and write access to the CIFS share.
- 2 In Windows, configure both the NetBackup Client Service and the NetBackup Remote Manager and Monitor Service to run under the same Windows user account that you created in step 1.

See your Windows documentation for details on how to configure the account for the services.

Disk storage units in storage lifecycle policies

Figure 11-2 is an example of how storage lifecycle policies can interact with volumes in a disk pool that a storage unit references.

Two backup policies are created as follows:


- A backup policy named Policy_gold has a gold classification. For storage, it is configured to use an SLP named Lifecycle_Gold, which has a gold data classification.
- A backup policy named Policy_silver has a silver classification. For storage, it is configured to use **Any Available**. That means it can use any available storage unit or any SLP that has a silver classification.

Two storage units are available to both backup policies as follows:

- DSU_1 is an operation in Lifecycle_Gold and references DiskPool_A.
- DSU_2 is not in an SLP and references DiskPool_A.

DiskPool_A contains three disk volumes. Both the gold and the silver images can be written to any disk volume in the pool.

Figure 11-2 Storage lifecycle policies and disk storage units referencing disk pools



See [“About storage lifecycle policies”](#) on page 601.

Maintain the available disk space on disk storage units

Disk storage units can be managed so that they do not become entirely full and cause backups to fail.

Create space for more images on a disk storage unit in the following ways:

- Add new disk space.
- Set the **High water mark** to a value that best works with the size of backup images in the environment.
See [“High water mark storage unit setting”](#) on page 561.

Maintain space on basic disk staging storage units in the following ways:

- Increase the frequency of the relocation schedule. Or, add resources so that all images can be copied to a final destination storage unit in a timely manner.
- Run the `nb_updatedssu` script.
Upon NetBackup installation or upgrade, the `nb_updatedssu` script runs. The script deletes the `.ds` files that were used in previous releases as pointers to relocated data. Relocated data is tracked differently in the current release and the `.ds` files are no longer necessary. Under some circumstances, a `.ds` file cannot be deleted upon installation or upgrade. In that case, run the script again:
On Windows: `install_path\NetBackup\bin\goodies\nb_updatedssu`
On UNIX: `/usr/opensv/netbackup/bin/goodies/nb_updatedssu`
- Determine the potential free space.
See [“Finding potential free space on a BasicDisk disk staging storage unit”](#) on page 584.
- Monitor disk space by enabling the **Check the capacity of disk storage units** host property.
This General Server host property determines how often NetBackup checks 6.0 disk storage units for available capacity. Subsequent releases use internal methods to monitor disk space more frequently.
See [“General server properties”](#) on page 106.

NDMP storage unit considerations

The NetBackup for NDMP license must be installed on the media server to use the hosts as storage units. Media Manager controls NDMP storage units but the units attach to NDMP hosts.

See [“About storage unit settings”](#) on page 559.

Create NDMP storage units for drives directly attached to NAS filers. Any drive that is attached to a NetBackup media server is considered a Media Manager storage unit, even if used for NDMP backups.

Note: Remote NDMP storage units may already be configured on a media server from a previous release. Upon upgrade of the media server, those storage units are automatically converted to Media Manager storage units.

See the [NetBackup for NDMP Administrator's Guide](#) for more information.

About storage unit settings

The following topics describe the settings that appear for all types of storage units. The settings are listed alphabetically. Each setting does not appear for each storage unit type.

Absolute pathname to directory or absolute pathname to volume setting for storage units

Absolute pathname to directory or **Absolute pathname to volume** is available for any storage unit that is not based on disk pools.

The setting specifies the absolute path to a file system or a volume available for backups to disk. Enter the path directly in the field, then click **Add**. Use any location on the disk, providing that sufficient space is available.

Use platform-specific file path separators (/ and \) and colon (:) within a drive specification.

The **Properties** button displays properties for the directory or volume.

See [“Properties option in the Change Storage Units dialog box”](#) on page 569.

Do not configure multiple BasicDisk storage units to use the same volume or file system. Not only do the storage units compete for space, but different **Low water marks** can cause unexpected behaviors.

If the BasicDisk storage unit is used as a disk staging storage unit, it is recommended to dedicate a disk partition or file system to it. Dedicating space allows the disk staging space management logic to operate successfully. Or, consider defining AdvancedDisk storage units, which use the disk pools that are composed of the disk volumes that are dedicated file systems for disk backup.

See [“NetBackup naming conventions”](#) on page 1054.

See [“Low water mark storage unit setting”](#) on page 562.

Directory can exist on the root file system or system disk setting for storage units

This setting allows the user to specify a directory on the root file system (UNIX) or on a system drive (Windows) in the **Absolute pathname to directory** field.

When this setting is checked, the directory is created automatically. If a storage unit is configured on c drive and this option is not checked, backups fail with status code 12 (file open failed).

Note: With this setting checked, the root file system or the system drive can fill up.

A job fails under the following conditions:

- If the setting is not checked, and if the directory already exists on the root file system or on a system drive.
- If the setting is not checked, and the requested directory is to be created on the root file system or on a system drive.

See [“Absolute pathname to directory or absolute pathname to volume setting for storage units”](#) on page 559.

Density storage unit setting

The **Storage device** selection determines the media **Density**. This setting appears for Media Manager and NDMP storage units only.

Disk type storage unit setting

The **Disk type** storage unit setting identifies the type of storage unit.

A disk storage unit can be one of the following types:

- **AdvancedDisk** (NetBackup Data Protection Optimization Option needed)
- **BasicDisk**
- **Cloud Storage** (NetBackup Data Protection Optimization Option needed)
- **OpenStorage (vendor name)** (NetBackup OpenStorage Disk Option needed)
- **PureDisk** (NetBackup Data Protection Optimization Option needed)

Note: The **null_stu** storage unit type is available only when Cohesity Technical Support uses the NullOST plug-in to identify and isolate data transfer bottlenecks. The **null_stu** storage unit type is used for troubleshooting purposes only. Do not select **null_stu** as a storage unit type because the data that is written to a null storage unit cannot be restored.

Enable block sharing storage unit setting

The **Enable block sharing** storage unit setting allows the sharing of data blocks that have not changed from one backup to the next. Sharing data blocks can significantly save disk space in the storage unit.

Enable multiplexing storage unit setting

The **Enable multiplexing** storage unit setting allows multiple backups to multiplex onto a single drive in a storage unit.

Caution: For MS-SQL-Server policies, do not enable multiplexing for a schedule that is also configured to backup with multiple stripes. Restores fail when multiplexing is enabled for a schedule that uses more than one stripe.

High water mark storage unit setting

The **High water mark** storage unit setting applies to **BasicDisk** storage units and to disk pools.

Note: **High water mark** does not apply to cloud storage disk pools. This value is derived from the storage capacity, which cannot be fetched from the cloud provider.

The **High water mark** setting (default 98%) is a threshold that triggers the following actions:

- When an individual disk volume of the underlying storage reaches the **High water mark**, NetBackup considers the volume full. NetBackup chooses a different volume in the underlying storage to write backup images to.
- When all volumes in the underlying storage reach the **High water mark**, the **BasicDisk** storage is considered full. NetBackup fails any backup jobs that are assigned to a storage unit in which the underlying storage is full. NetBackup also does not assign new jobs to a **BasicDisk** storage unit in which the underlying storage is full.

- NetBackup begins image cleanup when a volume reaches the **High water mark**; image cleanup expires the images that are no longer valid. NetBackup again assigns jobs to the storage unit when image cleanup reduces any disk volume's capacity to less than the **High water mark**.
If the storage unit is in a capacity-managed storage lifecycle policy, other factors affect image cleanup.
See [“Capacity managed retention type for SLP operations”](#) on page 624.

See [“Maximum concurrent jobs storage unit setting”](#) on page 563.

For more information, see the following guides:

- [NetBackup Deduplication Guide](#).
- [NetBackup Administrator's Guide, Volume II](#).

Low water mark storage unit setting

The **Low water mark** setting has no effect unless backups are written through a storage lifecycle policy, using the **Capacity managed** retention type. NetBackup copies expired images to a final destination storage unit to create space.

Note: **Low water mark** does not apply to cloud storage disk pools. This value is derived from the storage capacity, which cannot be fetched from the cloud provider.

Once the **High water mark** is reached, space is created on the disk storage unit until the **Low water mark** is met. The default setting is 80%.

See [“Capacity managed retention type for SLP operations”](#) on page 624.

The **Low water mark** setting cannot be greater than the **High water mark** setting.

For the disk storage units that reference disk pools, the **Low water mark** applies to the disk pool.

Note: Basic disk staging storage units may already be configured on a media server of a previous release. Upon upgrade, the disk storage units are set with the **Low water mark** at 100%. To make the best use of upgraded storage units, adjust the level.

For more information, see the following:

- [NetBackup Deduplication Guide](#).
- [NetBackup Administrator's Guide, Volume II](#).

Maximum concurrent write drives storage unit setting

The **Maximum concurrent write drives** storage unit setting specifies the number of tape drives that NetBackup can use at one time for jobs to this storage unit. The number of tape drives available is limited to the maximum number of tape drives in the storage device. If a job contains multiple copies, each copy applies toward the **Maximum concurrent write drives** count.

When selecting the value for **Maximum concurrent write drives**, use the following guidelines:

- Storage unit that contains only standalone tape drives
Specify a number that is less than or equal to the number of tape drives that are in the storage unit.
- Robot
Specify a number that is less than or equal to the number of tape drives that attach to the NetBackup media server for the storage unit.

Assume that you have two standalone drives of the same density and specify 1. Both tape drives are available to NetBackup but only one drive can be used for backups. The other tape drive is available for restores and other non-backup operations. (For example, to import, to verify, and to duplicate backups as source.)

Note: To specify a **Maximum concurrent write drives** setting of 0 disables the storage unit.

Maximum concurrent jobs storage unit setting

The **Maximum concurrent jobs** storage unit setting specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. The default setting is one job.

The maximum number of jobs that NetBackup can run concurrently is dependent on several factors and is not exclusively regulated by this setting. These factors include the following: the risk of reaching a disk full situation, the scheduling or capacity polling overhead, the media server I/O bandwidth, and various characteristics of the disk storage. No definitive method exists that can predict when a critical limit will be exceeded in a given system.

Note: To specify a **Maximum concurrent jobs** setting of 0 disables the storage unit.

For example, three backup jobs are ready to be sent to the storage unit and **Maximum concurrent jobs** is set to two. The first two jobs start while the third job

waits. If a job contains multiple copies, each copy applies toward the **Maximum concurrent jobs** count.

Note: Increase the **Maximum concurrent jobs** setting if the storage unit is used for catalog backups as well as non-catalog backups. Increase the setting to ensure that the catalog backup can proceed while regular backup activity occurs. Where disk pools are used, increase the setting if more than one server is in the storage unit.

The **Maximum concurrent jobs** setting uses and dependencies are as follows:

- Can be used to balance the load between disk storage units. A higher value (more concurrent jobs) means that the disk may be busier than if the value was set for fewer jobs.
The media server load balancing logic considers all storage units and all activity. A storage unit can indicate three media servers. If **Maximum concurrent jobs** is set to three and two of the media servers are busy or down, the third media server is assigned all three jobs.
- This setting depends on the available disk space and the server's ability to run multiple backup processes. Where disk pools are used, the setting also depends on the number of media servers in the storage unit.
If multiple storage units reference the same disk pool, the number of concurrent jobs that can access the pool is the sum of the **Maximum concurrent jobs** settings on all of the disk storage units. The setting applies to the storage unit and not to the disk pool. Therefore, the job load is automatically spread across the media servers that the storage unit configuration indicates.
- On Windows systems, even with multiple concurrent jobs, the time that is required for a job to complete depends on other factors:
 - The number of other jobs that are started at the same time.
 - The sequence in which the jobs were started.
 - The time that is required to complete each job.

See [“Impact when two disk storage units reference one disk pool”](#) on page 564.

Impact when two disk storage units reference one disk pool

[Figure 11-3](#) shows how the **Maximum concurrent jobs** settings are combined when two disk storage units share one disk pool.

In the example, DSU_1 is configured as follows:


- To use MediaServer_A
- To have a **Maximum concurrent jobs** setting of two
- To reference Disk_pool1

DSU_2 is configured as follows:

- To use MediaServer_B
- To have a **Maximum concurrent jobs** setting of three
- To reference Disk_pool1

Both storage units reference the same disk pool. Combined, the storage units have a **Maximum concurrent jobs** setting of five. However, only two jobs can run concurrently on MediaServer_A; three on MediaServer_B.

Figure 11-3 Impact when disk storage units use one disk pool but different media servers



If the storage units were configured to use both media servers, the media servers could run five concurrent jobs: two from DSU_1 and three from DSU_2.

See [“About storage unit settings”](#) on page 559.

Maximum streams per drive storage unit setting

The **Maximum streams per drive** storage unit setting determines the maximum number of concurrent, multiple client backups that NetBackup can multiplex onto a single drive. The range is from 2 to 32.

See [“Media multiplexing \(schedule attribute\)”](#) on page 756.

Media server storage unit setting

The **Media server** storage unit setting specifies one of the following:

- The NetBackup media server where the drives in the storage unit attach.
- The NetBackup media server that controls the disk storage unit.
- The NetBackup media servers that can write data to and read data from the disk pool.
- The NetBackup media servers that can move data to and from the disk pool.
- The NetBackup media servers that function as deduplication servers.

To make this storage unit available to any media server (default), select **Any Available**. NetBackup selects the media server dynamically at the time the policy is run.

Consider the following, depending on the type of storage.

Table 11-2 Media server setting details

Storage unit type	Considerations
BasicDisk	To configure a disk storage unit, select a single media server.
AdvancedDisk	The Media server setting specifies the NetBackup media servers that can write data to and read data from the disk pool. The media servers that are configured as storage servers appear in the media servers list. The disk storage must be directly attached to the media server that is configured as the storage server. NetBackup selects a media server when the policy runs.
Cloud storage	The Media server setting specifies the NetBackup media servers that can move data to or from the cloud vendor storage server. To allow any media server in the media server list to move data to the storage server, check Use Any Available Media Server. To restrict the media servers that can move data to the storage server, check Only Use The Following Media Servers. Then select the media servers that are allowed to move the data. Any media server in the list can receive data from the storage server; it does not have to be selected. A media server receives data for restore jobs and for storage monitoring purposes. Only the media servers on which storage server credentials are configured appear in the media servers list. If a server does not appear, verify that the software plug-in is installed and that login credentials are configured for that media server. NetBackup selects a media server when the policy runs.

Table 11-2 Media server setting details (*continued*)

Storage unit type	Considerations
NDMP	The Media server setting specifies the name of the media server that is to back up the NDMP host. Only those media servers that can talk to the specified NDMP storage device appear in the drop-down menu. An NDMP host can be authenticated on multiple media servers. Select Any Available to have NetBackup select the media server and storage unit at the time the policy is run.
OpenStorage	The Media server setting specifies the NetBackup media servers that can move data to or from the storage server. To allow any media server in the media server list to move data to the storage server, check Use Any Available Media Server. To restrict the media servers that can move data to the storage server, check Only Use The Following Media Servers. Then select the media servers that are allowed to move the data. Any media server in the list can receive data from the storage server; it does not have to be selected. A media server receives data for restore jobs and for storage monitoring purposes. Each media server that moves the data must meet the following requirements: ■ The vendor's software plug-in is installed. ■ The login credentials to the storage server are configured. Only the media servers on which storage server credentials are configured appear in the media servers list. If a server does not appear, verify that the software plug-in is installed and that login credentials are configured for that media server. Note: Run the <code>tpconfig</code> command line utility directly on the media server to configure and verify credentials. NetBackup selects a media server when the policy runs.
PureDisk (Media Server Deduplication Pool)	To allow any media server in the list to deduplicate data, select Use Any Available Media Server. To restrict the media servers that can deduplicate data, select Only Use The Following Media Servers. Then select the media servers that are allowed to deduplicate the data. Each media server must be configured as a deduplication media server. See the NetBackup Deduplication Guide.

See [“Use any available media server storage unit setting”](#) on page 574.

See [“Only use the following media servers storage unit setting”](#) on page 568.

NDMP host storage unit setting

The **NDMP host** storage unit setting specifies the NDMP tape server that is used to write data to tape. Select the host name from the drop-down menu or click **Add** to add a host.

On demand only storage unit setting

The **On demand only** storage unit setting specifies whether the storage unit is available exclusively on demand—that is, only when a policy or schedule is explicitly configured to use this storage unit. Uncheck **On demand only** to make the storage unit available to any policy or schedule.

Note: If **On demand only** is selected for all storage units, be sure to designate a specific storage unit for each policy or schedule. Otherwise, NetBackup is unable to find a storage unit to use.

Only use the following media servers storage unit setting

The **Only use the following media servers** storage unit setting restricts the media servers that are earmarked for storage. Check this setting and select the media servers that you want to use.

The following table describes the media server functionality for each type of storage.

Table 11-3 Media server functionality

Media server type	Functionality
AdvancedDisk storage media server	The media servers are both storage servers and data movers. The media servers that are configured as the storage servers and data movers appear in the media servers list.
Cloud storage media server	The media servers that are configured as data movers for the cloud storage implementation appear in the media server list. (For cloud storage, NetBackup media servers function as data movers.)

Table 11-3 Media server functionality (continued)

Media server type	Functionality
OpenStorage media server	The media servers that are configured as data movers for the OpenStorage implementation appear in the media server list. (For OpenStorage, NetBackup media servers function as data movers.) If a media server does not appear in the list, verify that the software plug-in is installed and that logon credentials are created. Each media server that accesses the storage must meet the following requirements: ■ The vendor's software plug-in is installed. ■ The login credentials to the storage server are configured.
PureDisk media server (Media Server Deduplication Pool)	The media servers function as deduplication servers. NetBackup deduplication must be configured.

See [“Use any available media server storage unit setting”](#) on page 574.

See [“Only use the following media servers storage unit setting”](#) on page 568.

Properties option in the Change Storage Units dialog box

Click **Properties** to display information about the volume or the disk pool, as follows:

Note: The following properties do not apply to cloud storage disk pools: **Available space**, **Capacity**, **High water mark**, **Low water mark**, **Raw size**, and **Usable size**.

All these values are derived from the storage capacity, which cannot be fetched from the cloud provider.

Table 11-4 Storage Units Properties

Property	Description
Available space	This value reflects the space that remains for storage on a disk storage unit. The following equation determines the available space: $\text{Available space} = \text{free space} + \text{potential free space} - \text{committed space}$ The <code>df</code> command may report a value for the available space that is slightly different from the actual free space value that appears as a result of the <code>nbdevquery</code> command: <pre>nbdevquery -listdv -stype server_type -dp disk_pool</pre> The available space that the <code>df</code> command lists does not include the space that the operating system reserves. Since NetBackup runs as <code>root</code>, the <code>nbdevquery</code> command includes the reserved space in the available space equation.
Capacity	The Capacity value reflects the total amount of space that the disk storage unit or pool contains, both used and unused.
Disk pool comments	Comments that are associated with the disk pool.
High water mark	The high water mark for the disk pool applies to both the individual disk volumes in the pool and the disk pool: ■ Individual volumes When a disk volume reaches the high water mark, new jobs are not assigned to the volume. This behavior happens for all disk types except BasicDisk staging storage units. The high water mark event triggers the deletion of images that have been relocated, attempting to bring the used capacity of the disk volume down to the low water mark ■ Disk pool When all volumes are at the high water mark, the disk pool is full. When a disk pool approaches the high water mark, NetBackup reduces the number of jobs that are allowed to write to the pool. NetBackup does not assign new jobs to a storage unit in which the disk pool is full. The default setting is 99%.
Low water mark	The low water mark for the disk pool. Once a disk volume fills to its high water mark, NetBackup attempts to delete enough relocated images to reduce the used capacity of the disk volume down to the low water mark. The low water mark setting cannot be greater than the high water mark setting. Note: The Low water mark setting has no effect unless backups are written through a storage lifecycle policy, using the capacity-managed retention type.
Name	The name of the disk pool.
Number of volumes	The number of disk volumes in the disk pool.

Table 11-4 Storage Units Properties (*continued*)

Property	Description
% full	The percentage of storage that is currently in use on the volume. The <code>df</code> command may report a percentage used (Use%) value that is different from the % full value. (See the preceding Available Storage topic for a description of why the values appear differently.)
Raw size	The raw, unformatted size of the storage in the disk pool.
Usable size	The amount of usable storage in the disk pools.

Reduce fragment size storage unit setting

The **Reduce fragment size** storage unit setting specifies the largest fragment size that NetBackup can create to store backups.

If an error occurs in a backup, the entire backup is discarded. The backup restarts from the beginning, not from the fragment where the error occurred. (An exception is for backups where checkpoint restart is enabled. In that case, fragments before and including the last checkpoint are retained; the fragments after the last checkpoint are discarded.)

Maximum fragment size for Media Manager storage units

The default maximum fragment size for a Media Manager storage unit is 1000 GB. To specify a maximum fragment size other than the default, check **Reduce fragment size**. Then enter a value from 50 megabytes to 1,048,575 megabytes.

Fragmenting multiplexed tape backups can expedite restores. Fragments allow NetBackup to skip to the specific fragment before searching for a file. Generally, NetBackup starts at the beginning of the multiplexed backup and reads tar headers until it finds the file.

Maximum fragment size for disk storage units

The default maximum fragment size for a disk storage unit is 524,288 megabytes. To specify a maximum fragment size other than the default, enter a value from 20 megabytes to 524,288 megabytes.

For a **Media Server Deduplication Pool**, you can enter a value from 20 megabytes to 51200 megabytes.

Backups to disk are usually fragmented to ensure that the backup does not exceed the maximum size that the file system allows.

The **Reduce fragment size** setting is intended primarily for storing large backup images on a disk type storage unit.

Note: OpenStorage vendors may have special requirements for the maximum fragment size. Consult the vendor's documentation for guidance.

Note: Basic disk staging units with different maximum fragment sizes may already be configured on a media server from a previous release. Upon upgrade, the disk storage units are not automatically increased to the new default of 524,288 megabytes. To make the best use of upgraded storage units, increase the fragment size on the upgraded storage units.

Robot number storage unit setting

The **Robot number** storage unit setting indicates the number of robots the storage unit contains. The **Storage device** selection determines the **Robot number**. It is the same robot number used in the Media Manager configuration.

Robot type storage unit setting

The **Robot type** storage unit setting indicates the type of robot (if any) that the storage unit contains. The **Storage device** setting determines the **Robot type**.
NetBackup robot types are described in a different topic.
See [“NetBackup robot types”](#) on page 418.

Select disk pool storage unit setting

The **Select disk pool** storage unit setting specifies disk pool for the storage unit.
The following table describes which disk pools appear in the drop-down list:

For AdvancedDisk	All NetBackup disk pools appear in the Disk pool list.
For cloud storage	Only the disk pools that the cloud storage vendor exposes appear in the list.
For OpenStorage	Only the disk pools that the OpenStorage vendor exposes appear in the list.
For PureDisk	The Media Server Deduplication Pools appear in the list.

Staging schedule option in Change Storage Units dialog

Click the **Staging Schedule** option to configure the relocation schedule for this storage unit. A schedule is what makes the disk storage unit a basic disk staging

storage unit. During the relocation schedule, the backup image is duplicated from the temporary staging area to the final destination storage unit.

See [“Disk Staging Schedule dialog box”](#) on page 585.

See [“Enable temporary staging area storage unit setting”](#) on page 573.

See [“About basic disk staging”](#) on page 577.

See [“About staging backups”](#) on page 576.

Storage device setting for storage units

The **Storage device** list contains all possible storage devices available. Storage units can be created for the listed devices only.

The **Storage device** selection determines the media **Density**. This setting appears for Media Manager and NDMP storage units only.

Storage unit name setting

The **Storage unit name** setting defines a unique name for the new storage unit. The name can describe the type of storage. The **Storage unit name** is the name used to specify a storage unit for policies and schedules.

The storage unit name cannot be changed after creation. The **Storage unit name** is inaccessible when changing settings for a storage unit.

See [“NetBackup naming conventions”](#) on page 1054.

Storage unit type setting

The **Storage unit type** setting specifies the type of storage that this storage unit uses, as follows:

Disk	See “Disk storage unit considerations” on page 554.
Media Manager	See “Tape storage unit considerations” on page 553.
NDMP	See “NDMP storage unit considerations” on page 558.

Enable temporary staging area storage unit setting

The **Enable temporary staging area** storage unit setting allows this storage unit to be used as a temporary staging area. Check **Enable Temporary Staging Area** and then configure the staging schedule.

See [“Staging schedule option in Change Storage Units dialog”](#) on page 572.

The Staging column in the **Storage units** details pane indicates whether or not the unit is used as a temporary staging area for basic disk staging. Not all columns display by default.

See [“About basic disk staging”](#) on page 577.

See [“Staging schedule option in Change Storage Units dialog”](#) on page 572.

Use any available media server storage unit setting

When checked, the **Use any available media server** storage unit setting allows any media server in the media server list to access the storage (default).

The following table describes the media server functionality for each type of storage.

Table 11-5 Media server functionality

Storage unit type	Functionality
AdvancedDisk storage media server	The media servers are both storage servers and data movers. The media servers that are configured as the storage servers and data movers appear in the media servers list.
Cloud storage media server	The media servers that are configured as data movers for the cloud storage implementation appear in the media server list. (For cloud storage, NetBackup media servers function as data movers.)
OpenStorage media server	The media servers that are configured as data movers for the OpenStorage implementation appear in the media server list. (For OpenStorage, NetBackup media servers function as data movers.) If a media server does not appear in the list, verify that the software plug-in is installed and that logon credentials are created. The following is required on each media server that accesses the storage: ■ The vendor's software plug-in is installed. ■ The login credentials to the storage server are configured.
PureDisk media server (Media Server Deduplication Pool)	The media servers function as deduplication servers. NetBackup deduplication must be configured.

Use WORM setting

The **Use WORM** option is enabled for storage units that are WORM capable. Select this option if you want the backup images on this storage unit to be immutable and indelible until the WORM Unlock Time.

Note: You must also select the **On Demand Only** option whenever the **Use WORM** option is selected.

WORM is the acronym for Write Once Read Many.

Overview of universal shares

The universal share feature provides data ingest into an existing NetBackup deduplication pool (MSDP) or a supported Cohesity appliance using an NFS or a CIFS (SMB) share.

Both universal shares and MSDP use deduplication and compression.

Space efficiency is achieved by storing this data directly into an existing NetBackup-based Media Server Deduplication Pool.

For more information about universal shares, see the [NetBackup Deduplication Guide](#).

Staging backups

This chapter includes the following topics:

- [About staging backups](#)
- [About basic disk staging](#)
- [Creating a basic disk staging storage unit](#)
- [Configuring multiple copies in a relocation schedule](#)
- [Disk staging storage unit size and capacity](#)
- [Finding potential free space on a BasicDisk disk staging storage unit](#)
- [Disk Staging Schedule dialog box](#)
- [Initiating a relocation schedule manually](#)

About staging backups

In the staged backups process, NetBackup writes a backup to a storage unit and then duplicates it to a second storage unit. Eligible backups are deleted on the initial storage unit when space is needed for more backups.

This two-stage process allows a NetBackup environment to leverage the advantages of disk-based backups for recovery in the short term.

Staging also meets the following objectives:

- Allows for faster restores from disk.
- Allows the backups to run when tape drives are scarce.
- Allows the data to be streamed to tape without image multiplexing.

NetBackup offers the following methods for staging backups.

Table 12-1 Methods for staging backups

Staging method	Description
Basic disk staging	Basic disk staging consists of two stages. First, data is stored on the initial storage unit (disk staging storage unit). Then, per a configurable relocation schedule, data is copied to the final location. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed. See “About basic disk staging” on page 577. The following storage unit types are available for basic disk staging: BasicDisk and tape.
Staging using the Storage lifecycle policies utility	Staged backups that are configured within the Storage lifecycle policies utility also consist of two stages. Data on the staging storage unit is copied to a final destination. However, the data is not copied per a specific schedule. Instead, the administrator can configure the data to remain on the storage unit until either a fixed retention period is met, or until the disk needs additional space, or until the data is duplicated to the final location. No BasicDisk or disk staging storage unit can be used in an SLP. See “About storage lifecycle policies” on page 601.

About basic disk staging

Basic disk staging is conducted in the following stages.


Table 12-2 Basic disk staging

Stage	Description
Stage I	Clients are backed up by a policy. The Policy storage selection in the policy indicates a storage unit that has a relocation schedule configured. The schedule is configured in the staging schedule settings.
Stage II	Images are copied from the Stage I disk staging storage unit to the Stage II storage unit. The relocation schedule on the disk staging storage unit determines when the images are copied to the final destination. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.

The image continues to exist on both the disk staging storage unit and the final destination storage units until the image expires or until space is needed on the disk staging storage unit.

[Figure 12-1](#) shows the stages in basic disk staging.

Figure 12-1 Stage I and II of basic disk staging



When the relocation schedule runs, NetBackup creates a data management job. The job looks for any data that can be copied from the disk staging storage unit to the final destination. The job details in the Activity monitor identify the job as one associated with basic disk staging. The jobs list displays Disk Staging in the job's **Data movement** field.

When NetBackup detects a disk staging storage unit that is full, it pauses the backup. Then, NetBackup finds the oldest images on the storage unit that successfully copied onto the final destination. NetBackup expires the images on the disk staging storage unit to create space.

Note: The basic disk staging method does not support backup images that span disk storage units.

To avoid spanning storage units, do not use Checkpoint restart on a backup policy that writes to a storage unit group that contains multiple disk staging storage units.

See [“Take checkpoints every __ minutes \(policy attribute\)”](#) on page 674.

Creating a basic disk staging storage unit

When you configure a BasicDisk storage unit with disk staging, data is stored on the initial storage unit (disk staging storage unit). Then with a configurable relocation schedule, data is copied to the final location. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.

To create a BasicDisk storage unit with disk staging

- 1 Open the NetBackup web UI.
- 2 Select **Storage > Storage units**.
- 3 Select **Add**.
- 4 Select **BasicDisk**. Then click **Start**.

5 Select the properties for the storage unit.

Type a **Name** for the storage unit.

See [“Storage unit name setting”](#) on page 573.

Enter the number of **Maximum concurrent jobs** that are allowed to write to this storage unit at one time.

See [“Maximum concurrent jobs storage unit setting”](#) on page 563.

Enter a **High water mark** value.

The high water mark works differently for the BasicDisk disk type. NetBackup assigns new jobs to a BasicDisk disk staging storage unit, even if it is over the indicated high water mark. For BasicDisk, the high water mark is used to prompt the deletion of images that have been relocated.

Note: The **Low water mark** setting does not apply to disk staging storage units.

6 Select **Next**.

7 For the staging schedule, select the option **Enable temporary staging area**.

8 In the **Staging schedule** section, then select **Add**.

The schedule name defaults to the storage unit name.

Configure the schedule settings.

See [“Disk Staging Schedule dialog box”](#) on page 585.

9 Click **Save** to save the disk staging schedule.

10 Click **Next**.

11 Select a media server.

See [“Media server storage unit setting”](#) on page 565.

12 Browse or specify the absolute path to the directory to use for storage.

See [“Absolute pathname to directory or absolute pathname to volume setting for storage units”](#) on page 559.

13 Select whether this directory can reside on the root file system or system disk.

See [“Directory can exist on the root file system or system disk setting for storage units”](#) on page 560.

- 14** Click **Next**.
- 15** Review the settings for the storage unit and then click **Save**.

Creating a schedule for a BasicDisk staging storage unit

When you configure a BasicDisk storage unit with disk staging, data is stored on the initial storage unit (disk staging storage unit). Then, per a configurable relocation schedule, data is copied to the final location. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.

To create a schedule for a BasicDisk staging storage unit

- 1** Open the NetBackup web UI.
- 2** On the left, click **Storage > Storage units**. Then click on the BasicDisk storage unit to which you want to add a schedule.
- 3** To the right of **Staging schedule**, click **Edit**.
- 4** Select **Enable a temporary staging area**.
The schedule name defaults to the storage unit name.
- 5** Configure the setting **Priority of relocation job started from this schedule**. This setting controls the priority that relocation jobs have compared to other types of jobs.
See [“Disk Staging Schedule dialog box”](#) on page 585.
- 6** Select whether to create multiple copies. When the **Multiple copies** attribute is selected, NetBackup can create up to four copies of a backup simultaneously.
See [“Multiple copies \(schedule attribute\)”](#) on page 747.
- 7** Select a **Final destination storage unit** to contain the images from this storage unit upon relocation.
- 8** Select a **Final destination volume pool** to contain the images from this storage unit upon relocation.
- 9** Select a **Final destination media ownership** setting that controls the media owner that owns the images from this storage unit upon relocation.
- 10** Choose the **Schedule type**.
- 11** Select whether to **Use alternate read server** for the images from this storage unit upon relocation.
- 12** Click **Save** to save the disk staging schedule.

Configuring multiple copies in a relocation schedule

To configure a relocation schedule for basic disk staging to create multiple copies, use the following procedure.

To configure a relocation schedule for basic disk staging to create multiple copies

- 1 Open the NetBackup web UI.
- 2 Open the storage unit that you want to edit. Then edit the **Staging schedule**.
- 3 On the **Attributes** tab, select **Multiple copies**.
- 4 In the **Copies** field, specify the number of copies to create simultaneously. The number must be between 1 and 4.
- 5 Specify a priority in the field **Priority of relocation jobs started from this schedule** (0 to 99999).
- 6 For each copy you want to create, select the copy settings. **Copy 1** is the primary copy. If **Copy 1** fails, the first successful copy is the primary copy.
See [“Configure Multiple Copies dialog box”](#) on page 748.
See [“About configuring multiple copies”](#) on page 747.
- 7 Specify the storage unit where each copy is stored. If a Media Manager storage unit has multiple drives, it can be used for both the source and the destination.
- 8 Specify the volume pool where each copy is stored.
- 9 Select one of the following from the **If this copy fails** list:

Continue	Continues making the remaining copies. Note: Note: If Take checkpoints every __ minutes is selected for this policy, only the last failed copy that contains a checkpoint can be resumed. See “Take checkpoints every __ minutes (policy attribute)” on page 674.
Fail all copies	Fails the entire job.

- 10** For tape media, specify who should own the media onto which NetBackup writes the images:

Any	NetBackup selects the media owner, either a media server or server group.
None	Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All media server groups that are configured in the NetBackup environment appear in the list.

These settings do not affect images that reside on disk. One media server does not own the images that reside on shared disks. Any media server with access to the shared pool of disk can access the images.

- 11** Click **Add** or **Save**.

Disk staging storage unit size and capacity

To take advantage of basic disk staging requires that the NetBackup administrator understand the life expectancy of the image on the Stage I storage unit.

The size and use of the file system of the Stage I storage unit directly affects the life expectancy of the image before it is copied to the Stage II storage unit. It is recommended a dedicated file system for each disk staging storage unit.

Consider the following example: A NetBackup administrator wants incremental backups to be available on disk for one week.

Incremental backups are done Monday through Saturday, with full backups done on Sunday. The full backups are sent directly to tape and do not use basic disk staging.

Each night's total incremental backups are sent to a disk staging storage unit and average from 300 MB to 500 MB. Occasionally a backup is 700 MB. Each following day the relocation schedule runs on the disk staging storage unit and copies the previous night's incremental backups to the final destination, a Media Manager (tape) storage unit.

The following items give more information about determining disk size for a basic disk staging storage unit.

Minimum disk size

The minimum disk size is the smallest size that is required for the successful operation of the disk staging logic.

The minimum size must be greater than or equal to the largest combined size of the backups that are placed on the storage unit between runs of the disk staging schedule. (In our example, the disk images remain on the disk for one week.)

In this example, the relocation schedule runs nightly, and the largest nightly backup is 700 MB. It is recommended that you double this value to allow for any problems that may occur when the relocation schedule runs. To double the value gives the administrator an extra schedule cycle (one day) to correct any problems.

To determine the minimum size for the storage unit in this example, use the following formula:

Minimum size = Max data per cycle × (1 cycle + 1 cycle for safety)

For example: 1.4 GB = 700 MB × (1+1)

Average disk size

The average disk size represents a good compromise between the minimum and the maximum sizes.

In this example, the average nightly backup is 400 MB and the NetBackup administrator wants to keep the images for one week.

To determine the average size for the storage unit in this example, use the following formula:

Average size = Average data per cycle × (number of cycles to keep data + 1 cycle for safety)

2.8 GB = 400 MB × (6 + 1)

Maximum disk size

The maximum disk size is the recommended size needed to accommodate a certain level of service. In this example, the level of service is that disk images remain on disk for one week.

To determine the maximum size for the storage unit in this example, use the following formula:

Maximum size = Max data per cycle × (# of cycles to keep data + 1 cycle for safety)

For example: 4.9 GB = 700 MB × (6 + 1)

Finding potential free space on a BasicDisk disk staging storage unit

Potential free space is the amount of space on a disk staging storage unit that NetBackup could free if extra space on the volume is needed. The space is the total size of the images that are eligible for expiration plus the images ready to be deleted on the volume.

To find the potential free space on a BasicDisk storage unit, use the `bpstulist` and the `nbdevquery` commands as follows:

- Run `bpstulist -label` to find the disk pool name.
Note that the name of the storage unit and disk pools are case-sensitive. In the case of BasicDisk storage units, the name of the disk pool is the same as the name of the BasicDisk storage unit. In the following example, the name of the storage unit is *NameBasic*:

```
bpstulist -label basic
NameBasic 0 server1 0 -1 -1 1 0 "C:\\" 1 1 524288 *NULL* 0 1 0 98 80 0 NameBasic server1
```

- Run the `nbdevquery` command to display the status for the disk pool, including the potential free space.
Use the following options, where:

<code>-stype server_type</code>	Specifies the vendor-specific string that identifies the storage server type. For a BasicDisk storage unit, enter <code>BasicDisk</code> .
<code>-dp</code>	Specifies the disk pool name. For a basic disk type, the disk pool name is the name of the BasicDisk storage unit.

So the complete command might look like the following.

```
nbdevquery -listdv -stype BasicDisk -dp NameBasic -D
```

The value is listed as `potential_free_space`.

```
Disk Volume Dump
name           : <Internal_16>
id             : <C:\>
diskpool       : <NameBasic::server1::BasicDisk>
disk_media_id  : <@aaaaf>
total_capacity : 0
free_space     : 0
```

```

potential_free_space: 0
committed_space      : 0
precommitted_space   : 0
nbu_state             : 2
sts_state             : 0
flags                 : 0x6
num_read_mounts       : 0
max_read_mounts       : 0
num_write_mounts      : 1
max_write_mounts      : 1
system_tag            : <Generic disk volume>

```

Disk Staging Schedule dialog box

Click **Staging Schedule** to display the **Disk Staging Schedule** dialog box. The dialog box is similar to the scheduling dialog box that appears when a policy is configured.

The schedule that is created for the disk staging storage unit is not listed under **Schedules** in the **NetBackup Administration Console** when the **Policies** utility is selected.

The following settings are available when you create a disk staging schedule.

Table 12-3 The Attributes tab settings

Attribute	Description
Name	The schedule Name defaults to the name of the storage unit.
Priority of relocation jobs started from this schedule	The Priority of relocation jobs started from this schedule field indicates the priority that NetBackup assigns to relocation jobs for this policy. Range: 0 to 99999 (highest priority). The default value that is displayed is the value that is set in the Default job priorities host properties for the Staging job type. See “Default job priorities properties” on page 85.
Multiple copies	Creates multiple copies of backups. NetBackup can create up to four copies of a backup simultaneously. When this setting is enabled, Final destination volume pool and Final destination media ownership are disabled. See “Multiple copies (schedule attribute)” on page 747.

Table 12-3 The Attributes tab settings (*continued*)

Attribute	Description
Final destination storage unit	If the schedule is a relocation schedule, a Final destination storage unit must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A Final destination storage unit is the name of the storage unit where the images reside after a relocation job copies them. To copy images to tape, NetBackup uses all of the drives available in the Final destination storage unit. However, the Maximum concurrent write drives setting for that storage unit must be set to reflect the number of drives. The setting determines how many duplication jobs can be launched to handle the relocation job. NetBackup continues to free space until the Low water mark is reached. See “Low water mark storage unit setting” on page 562. See “Maximum concurrent write drives storage unit setting” on page 563. See “About staging backups” on page 576.
Final destination volume pool	If the schedule is a relocation schedule, a Final destination volume pool must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A Final destination volume pool is the volume pool where images are swept from the volume pool on the basic disk staging storage unit. See “About staging backups” on page 576. Note: The relocation schedule that was created for the basic disk staging storage unit is not listed under Schedules in the NetBackup Administration Console when the Policies utility is selected.
Final destination media owner	If the schedule is a relocation schedule, a Final destination media owner must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A Final destination media owner is the media owner where the images reside after a relocation job copies them. Specify one of the following: ■ Any lets NetBackup choose the media owner. NetBackup chooses a media server or a server group (if one is configured). ■ None specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ A server group. A server group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the Final destination media owner drop-down list.

Table 12-3 The Attributes tab settings (*continued*)

Attribute	Description
Schedule type	Calendar See “Calendar (schedule attribute)” on page 744. Frequency See “Frequency (schedule attribute)” on page 744. If the backups that use a disk staging storage unit run more frequently than expected, compare the retention level 1 setting with the Frequency setting. Internally, NetBackup uses the retention level 1 setting for scheduling purposes with disk staging storage units. Make sure that the frequency period is set to make the backups occur more frequently than the retention level 1 setting indicates. (The default is two weeks.) For example, a frequency of one day and a retention level 1 of 2 weeks should work well. Retention levels are configured in the Retention periods host properties. See “Retention periods properties” on page 152.
Use alternate read server	An alternate read server is a server allowed to read a backup image originally written by a different media server. The path to the disk or directory must be identical for each media server that is to access the disk. If the backup image is on tape, the media servers must share the same tape library or the operator must find the media. If the backup image is on a robot that is not shared or a standalone drive, the media must be moved to the new location. An administrator must move the media, inventory the media in the new robot, and run <code>bpmedia -oldserver -newserver</code> or assign a failover media server. To avoid sending data over the network during duplication, specify an alternate read server that meets the following conditions: ■ Connected to the storage device that contains the original backups (the source volumes). ■ Connected to the storage device that contains the final destination storage units. If the final destination storage unit is not connected to the alternate read server, data is sent over the network.
Copies	Specify the number of copies to create simultaneously. Range: 1 to 4.
Priority of duplication job	Indicates the priority that NetBackup assigns to duplication jobs for this policy. Range: 0 to 99999 (highest priority).

Table 12-3 The Attributes tab settings (*continued*)

Attribute	Description
Copy #	For each copy you want to create, select the copy settings. Copy 1 is the primary copy. If Copy 1 fails, the first successful copy is the primary copy. Storage unit Specify the storage unit where each copy is stored. If a Media Manager storage unit has multiple drives, it can be used for both the source and the destination. Volume pool Specify the volume pool where each copy is stored. If this copy fails ■ Continue Continues making the remaining copies. Note: Note: If Take checkpoints every __ minutes is selected for this policy, only the last failed copy that contains a checkpoint can be resumed. See "Take checkpoints every __ minutes (policy attribute)" on page 674. ■ Fail all copies Fails the entire job. Media owner For tape media, specify who should own the media onto which NetBackup writes the images. These settings do not affect any images that reside on disk. One media server does not own the images that reside on shared disks. Any media server with access to the shared pool of disk can access the images. ■ Any NetBackup selects the media owner, either a media server or server group. ■ None Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ A server group Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All media server groups that are configured in the NetBackup environment appear in the list.

Initiating a relocation schedule manually

A relocation schedule may be started manually to copy images to the final destination before the schedule is due to run.

To initiate a relocation schedule

- 1** In the **NetBackup Administration Console**, select **NetBackup Management > Storage > Storage Units**.
- 2** In the right pane, select a basic disk staging storage unit.
- 3** Select **Actions > Manual Relocation** or **Manual Relocation to Final Destination** to initiate the schedule.

If the relocation schedule finds data that can be copied, NetBackup creates a job to copy the data to the final destination storage unit.

The image then exists on both storage units until the disk staging (Stage I) storage unit becomes full and the oldest images are deleted.

See [“Maintain the available disk space on disk storage units”](#) on page 558.

Configuring storage unit groups

This chapter includes the following topics:

- [About storage unit groups](#)
- [Creating storage unit groups for backups](#)
- [Creating storage unit groups for snapshots](#)
- [Deleting a storage unit group](#)
- [Storage unit selection criteria within a group](#)
- [About disk spanning within storage unit groups](#)

About storage unit groups

Storage unit groups let you identify specific storage units as a group. You can specify a storage unit group name as the storage for a policy in the same way that you specify individual storage units. When you specify a storage unit group, the policy directs backups or snapshots only to those storage units in the designated group.

Storage unit groups can be one of the following types:

- **Backup storage unit groups**
A backup storage unit group contains only the storage units that can contain backups. Furthermore, for **Media Server Deduplication Pool** and third-party disk appliance (OpenStorage) storage, all storage units in the group must be of the same type of storage.
See [“Creating storage unit groups for backups”](#) on page 591.

- Snapshot storage unit groups
A snapshot storage unit group contains only the storage units that can contain snapshots. All storage units in the group must have the same **Disk type** selected. See [“Creating storage unit groups for snapshots”](#) on page 593.

NetBackup does not support storage unit groups for the following use cases:

- As a target for optimized duplication.
If you use a storage unit group as a destination for optimized duplication of deduplicated data, NetBackup uses regular duplication.
- As a source of or a target for Auto Image Replication.
- As a target for optimized synthetic backups.
If NetBackup cannot produce the optimized synthetic backup, NetBackup creates the more data-movement intensive synthetic backup.
- As a target for OpenStorage direct-to-tape operations.
See the [NetBackup OpenStorage Solutions Guide for Disk](#).

Creating storage unit groups for backups

The following procedure describes how to create a storage unit group that consists of the storage units that can contain backups.

To create a storage unit group

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Storage**.
- 2 Right-click **Storage Unit Groups** and select **New Storage Unit Group**.
- 3 Enter a storage unit group name for the new storage unit group. The storage unit group name is case-sensitive.
See [“NetBackup naming conventions”](#) on page 1054.
- 4 For the storage unit group to contain backups, select **Backup** in the drop-down menu.
- 5 Add backup storage units to or remove backup storage units from the group:
 - To add storage units to the group, select the storage units from the **Storage units not in the group** list and click **Add**.
 - To remove storage units from the group, select the storage units from the **Storage units in group** list and click **Remove**.
 - To change the priority of a storage unit, select the storage unit and click **Move Up** or **Move Down**. The units at the top of the list have the highest priority in the group.

Note: For **Media Server Deduplication Pool** and third-party disk appliance (OpenStorage) storage, all storage units in the group must be of the same type of storage.

6 Choose how storage units are selected within the group:

- **Prioritized.** Choose the first storage unit in the list that is not busy, down, or out of media.
- **Failover.** Choose the first storage unit in the list that is not down or out of media.

It is recommended that you select **Failover** for the following storage types: **AdvancedDisk**, **Media Server Deduplication Pool**, and **OpenStorage (VendorName)**.

- **Round Robin.** Choose the least recently selected storage unit in the list.
- **Media server load balancing.** Choose a storage unit based on a capacity-managed approach.

Media server load balancing is recommended for disk staging storage units within a storage unit group.

See [“Media server load balancing”](#) on page 596.

See [“Storage unit selection criteria within a group”](#) on page 595.


One exception to the selection criteria is in the case of a client that is also a media server with locally connected storage units.

See [“Exception to the storage unit selection criteria”](#) on page 598.

7 Click **OK**.

See [“About storage unit groups”](#) on page 590.

Figure 13-1 Backup storage unit group configuration dialog box



Creating storage unit groups for snapshots

Creating storage unit groups is optional. A snapshot storage unit group must be comprised of storage units that have matching properties.

The following procedure describes how to create a storage unit group that consists of the storage units that can contain snapshots.

To create a snapshot storage unit group

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management** > **Storage**.
- 2 Right-click **Storage Unit Groups** and select **New Storage Unit Group**.
- 3 Enter a storage unit group name for the new storage unit group. The storage unit group name is case-sensitive.
 See [“NetBackup naming conventions”](#) on page 1054.
- 4 For the storage unit group to contain snapshots, select **Snapshot** in the drop-down menu.
- 5 A storage unit group can contain only those storage units that share similar properties. NetBackup filters the storage units for selection so that dissimilar storage units are not combined in one storage unit group.

Note: The properties of the underlying storage units are read-only. You cannot change the storage unit properties from this dialog box.

Select one or more properties to filter the storage units in the list. Only those storage units that have the selected properties are displayed. For example, select **Replication source** and **Replication target** to display only those storage units that are configured to act as both replication sources and replication targets.


Filter the storage units on the following properties:

- **Primary**
 Enable **Primary** to display the storage units that can contain the initial snapshot of primary data.
- **Replication source**
 Enable **Replication source** to display the storage units that can serve as a source for a replicated snapshot.
- **Replication target**
 Enable **Replication target** to display the storage units that can receive replicated snapshots from a replication source.
 - **Mirror**
 Optionally, enable **Mirror** to display the storage units that can serve as a mirrored replication target. (For example, NetApp SnapMirror.)
 - **Independent**
 Optionally, enable **Independent** to display the storage units that can act as either a **Mirror** replication target (SnapMirror) or a non-mirror replication target (SnapVault).

For more information about snapshot properties, see the [NetBackup Replication Director Solutions Guide](#).

- 6 Add or remove storage units from the group:
 - To add storage units to the group, select the storage units from the **Available storage units** list and click **Add**.
 - To remove storage units from the group, select the storage units from the **Storage units to be used in the group** list and click **Remove**.
 - To change the priority of a storage unit, select the storage unit and click **Move Up** or **Move Down**. The units at the top of the list have the highest priority in the group.
- 7 Click **OK** to save and close the dialog box.

Figure 13-2 Snapshot storage unit group configuration dialog box



Deleting a storage unit group

The following procedure describes how to delete a storage unit group.

To delete a storage unit group

- 1 In the **NetBackup Administration Console**, select **NetBackup Management > Storage > Storage Unit Groups**.
- 2 In the right pane, from the list of storage unit groups, select the storage unit group you want to delete. Hold down the **Control** or **Shift** key to select multiple storage units.
- 3 Select **Edit > Delete**.
- 4 Click **OK**.

Storage unit selection criteria within a group

The storage unit selection criteria determines the order in which storage units are selected within a storage unit group.

The only difference between the selection criteria options is the order in which the storage units are selected.

Choose from one of the following selection criteria.

Selection	Description
Prioritized	If the Prioritized option is selected, NetBackup chooses the next available storage unit in the list. Prioritized is the default selection. If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.
Failover	If the Failover option is selected, when a job must queue for a storage unit, the job queues rather than try another storage unit in the group.
Round robin	If the Round robin option is selected, NetBackup chooses the least recently selected storage unit in the list as each new job is started. If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.
Media server load balancing	If the Media server load balancing option is selected, NetBackup selects a storage unit based on a capacity-managed approach. In this way, NetBackup avoids sending jobs to busy media servers. If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available. See “Media server load balancing” on page 596.

A queue can form for a storage unit if the storage unit is unavailable.

The following are some reasons why a storage unit can be considered unavailable:

- The storage unit is busy.
- The storage unit is down.
- The storage unit is out of media.
- The storage unit has no available space.
- The storage unit has reached the **Maximum concurrent jobs** setting.

See [“Maximum concurrent jobs storage unit setting”](#) on page 563.

See [“Exception to the storage unit selection criteria”](#) on page 598.

Media server load balancing

The **Media server load balancing** option indicates that NetBackup select a storage unit based on a capacity-managed approach. In this way, NetBackup avoids sending jobs to busy media servers.

If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.

The selection is based on the following factors:

- The rank of the media server.
 NetBackup considers the number of processes that are running on each CPU along with the memory thresholds on each server to determine the rank of a media server. If the free memory drops below a determined threshold, or if the number of running processes per CPU rises over a determined threshold, then the overall rank of the media server drops.
- The number of jobs on the media server.
 NetBackup considers the number of scheduled jobs on each media server.
- Whether the media server has enough disk space to accommodate the estimated size of the image. (Physical and virtual tapes ignore this requirement.)
 NetBackup estimates the size of any of the new or any current jobs on each media server. It then determines whether the jobs fit on a given volume.
 NetBackup estimates the amount of space that the job may require, based on previous backup history. If no history is available, the high water mark for the storage unit serves as a guide.

Media server load balancing cannot be selected for a storage unit group that includes a BasicDisk storage unit. Also, a BasicDisk storage unit cannot be included in an existing storage unit group with **Media server load balancing** enabled.

Note: It is recommended that you select **Media server load balancing** for disk staging storage units within a storage unit group.

The following methods also work to distribute the backup workload:

Adjust the backup load on a media server.

- Change the **Limit jobs per policy** policy attribute for one or more of the policies that are sent to a media server. Specifying a lower limit reduces the workload on a media server on a specific network segment.
 See [“Limit jobs per policy \(policy attribute\)”](#) on page 678.
- Reconfigure policies or schedules to use storage units on other media servers.
- Consider changing the **Bandwidth** host properties on one or more clients.
 See [“Storage unit selection criteria within a group”](#) on page 595.

Distribute the backup load on media servers during peak periods.

Reconfigure policy schedules so that they write backups to storage units on the media servers that can handle the load (assuming that master servers and media servers are on separate hosts).

Adjust the backup load on the client.	Change the Maximum jobs per client global attribute. For example, raising the Maximum jobs per client limit increases the number of concurrent jobs that any one client can process and therefore increases the load. See “Storage unit selection criteria within a group” on page 595.
Reduce the time that is needed to back up clients.	Increase the number of jobs that clients can perform concurrently, or use multiplexing. Another possibility is to increase the number of jobs that the media server can perform concurrently for the policies that back up the clients.
Give preference to a policy.	Increase the Limit jobs per policy attribute for the preferred policy relative to other policies. Or, increase the priority for the policy. See “Limit jobs per policy (policy attribute)” on page 678.
Adjust the load between fast and slow networks.	Increase the Limit jobs per policy and Maximum jobs per client for policies and clients in a faster network. Decrease these numbers for slower networks. Another solution is to use the NetBackup capability to limit bandwidth. See “Limit jobs per policy (policy attribute)” on page 678. See “Storage unit selection criteria within a group” on page 595.
Maximize the use of devices.	Use multiplexing. Allow as many concurrent jobs per storage unit, policy, and client as possible without causing server, client, or network performance problems.
Prevent backups from monopolizing tape devices.	■ Place some drives in a down state or limit the number that are used concurrently in a specific storage unit. For example, if there are four drives in a robot, allow only two to be used concurrently. ■ Do not place all devices under Media Manager control.

Exception to the storage unit selection criteria

The only exception to the storage unit selection criteria order is in the case of a client that is also a media server with locally connected storage units. The locally available storage units take precedence over the defined sequence of storage units in the group.

You may have set up a storage unit to be **On demand only**. If the unit is in a storage unit group that a policy requires, the **On demand only** option is satisfied and the device is used.

See [“On demand only storage unit setting”](#) on page 568.

See [“Storage unit selection criteria within a group”](#) on page 595.

About disk spanning within storage unit groups

A backup may span storage units if a disk full condition is detected. Backups can span from one BasicDisk storage unit to another BasicDisk storage unit if the storage units are in the same storage unit group. The storage units must also share the same media server.

See [“Storage unit selection criteria within a group”](#) on page 595.

Configuring storage lifecycle policies (SLPs)

- [Chapter 14. Configuring storage lifecycle policies](#)
- [Chapter 15. Storage operations](#)
- [Chapter 16. Retention types for SLP operations](#)
- [Chapter 17. Storage lifecycle policy options](#)
- [Chapter 18. Using a storage lifecycle policy to create multiple copies](#)
- [Chapter 19. Storage lifecycle policy versions](#)

Configuring storage lifecycle policies

This chapter includes the following topics:

- [About storage lifecycle policies](#)
- [Creating a storage lifecycle policy](#)
- [Deleting a storage lifecycle policy](#)
- [Lifecycle operation administration using the nbstlutil command](#)

About storage lifecycle policies

A storage lifecycle policy (SLP) is a storage plan for a set of backups. An SLP is configured within the **Storage lifecycle policies** utility. It contains instructions in the form of storage operations, to be applied to the data that is backed up by a backup policy. Operations are added to the SLP that determine how the data is stored, copied, replicated, and retained. NetBackup retries the copies as necessary to ensure that all copies are created.

SLPs offer the opportunity for users to assign a classification to the data at the policy level. A data classification represents a set of backup requirements, which makes it easier to configure backups for data with different requirements. For example, email data and financial data.

SLPs can be set up to provide staged backup behavior. They simplify data management by applying a prescribed behavior to all the backup images that are included in the SLP. This process allows the NetBackup administrator to leverage the advantages of disk-based backups in the near term. It also preserves the advantages of tape-based backups for long-term storage.

The **SLP settings** host properties allow administrators to customize how SLPs are maintained and how SLP jobs run.

Best practice information about SLPs appears in the following document:

https://www.veritas.com/content/support/en_US/article.100009913

Creating a storage lifecycle policy

A storage lifecycle policy (SLP) is a storage plan for a set of backups. The operations in an SLP are the backup instructions for the data. Use the following procedure to create an SLP that contains multiple storage operations.

To add a storage operation to a storage lifecycle policy

- 1 On the left, select **Storage > Storage lifecycle policies**.
- 2 Click **Add**.
- 3 Enter the Storage lifecycle policy name.
See “[Storage lifecycle policy settings](#)” on page 629.
- 4 Add one or more operations to the SLP. The operations are the instructions for the SLP to follow and apply to the data that is specified in the backup policy.
If this is the first operation added to the SLP, click **Add**.
To add a child operation, select an operation and then click **Add child**.
- 5 Select an **Operation** type. If you're creating a child operation, the SLP displays only those operations that are valid based on the parent operation that you selected.
See “[Operation types in a storage lifecycle policy](#)” on page 607.
- 6 Configure the properties for the operation.
- 7 The **Window** tab displays for the following operation types: **Backup From Snapshot**, **Duplication**, **Import**, **Index From Snapshot**, and **Replication**. If you'd like to control when the secondary operation runs, create a window for the operation.
- 8 On the **Properties** tab, go to and expand **Advanced**. Choose if NetBackup should process active images after the window closes.
- 9 Select **Create** to create the operation.
- 10 Add additional operations to the SLP as needed. (See step 4.)
- 11 Change the hierarchy of the operations in the SLP if necessary.

- 12 Select **Create** to create the SLP. NetBackup validates the SLP when it is first created and whenever it is changed.
- 13 Configure a backup policy and select a storage lifecycle policy as the **Policy storage**.

See [“Creating a backup policy”](#) on page 657.

Modifying the hierarchy of operations in a storage lifecycle policy

In some cases, the hierarchy of operations in an SLP can be modified. Use the arrows to move it in the hierarchy.

Note: It is not recommended that users modify automanaged storage lifecycle policies. If a user begins to modify an automanaged SLP, a dialog appears that warns users about the possible consequences.

See [“About automanaged policies or storage lifecycle policies”](#) on page 660.

The order of the operations at the time that the SLP is saved may differ from the next time the SLP is opened. NetBackup reorders the operations while it stores them in the catalog configuration file. How the hierarchy works is not changed, however, and the parent-child relationships are preserved.

Modify the order of the operations in the list if needed. Changing the order does not change the source of the selected operation. Changing the order does move the children of an operation and preserves their relationship with the selected operation.

- Select **Edit > Up** or **Edit > Down** to move an operation up or down in the hierarchy.
- Select **Edit > Promote** or **Edit > Demote** to make the selected operation a sibling of the parent operation or child of a sibling operation. You cannot promote or demote **Backup** or **Snapshot** operations.

Deleting a storage lifecycle policy

To delete a storage lifecycle policy, use the following procedure. Note that to delete an SLP deletes all versions of the SLP.

Note: It is not recommended that users modify or delete automanaged storage lifecycle policies. If a user begins to modify or delete an automanaged SLP, a dialog appears that warns users about the possible consequences.

See [“About automanaged policies or storage lifecycle policies”](#) on page 660.

To delete a storage lifecycle policy

- 1 Remove the SLP from all backup policies to prevent new backup jobs from writing to the SLP.
- 2 Wait for all in-process backup jobs to the SLP to complete, or cancel the jobs using the **Activity monitor** or the command line.
- 3 To prevent new jobs or cancel any existing duplication jobs from writing to the SLP, run the following command:

```
nbstlutil cancel -lifecycle name
```

- 4 Use the **Activity Monitor** to cancel in-process jobs that use the SLP.
- 5 Once all of the operations are complete, delete the SLP using one of the following methods:

- The **NetBackup Administration Console**
 - Expand **Storage > Storage Lifecycle Policies**.
 - Select the SLP name.
 - Select **Edit > Delete**.
 - In the **Delete Storage Lifecycle Policies** dialog box, select the SLP name and click **OK**.
- The `nbstl` command

```
nbstl storage_lifecycle_name -delete
```

If the administrator tries to delete an SLP with active images, status code 1519 appears (Images are in process). Wait several minutes and try to delete the SLP again until the error no longer appears.

Note: If orphaned images are detected due to a system error, NetBackup logs the fact that the images exist and alerts the administrator to address the situation.

Lifecycle operation administration using the nbstlutil command

The NetBackup storage lifecycle policy utility command (`nbstlutil`) gives administrators the ability to intervene between pending SLP operations. Specifically, the `nbstlutil` command can be used to cancel, inactivate, or activate the processing of existing SLP-managed images.

`nbstlutil` cannot affect the jobs that are currently running or queued. Use the **Activity Monitor** to intervene in the jobs that are running or queued.

Table 14-1 nbstlutil details

nbstlutil information	Details
Where to find	The command is found in the following location: On Windows: <pre>install_path\NetBackup\bin\admincmd\nbstlutil</pre> On UNIX: <pre>/usr/opensv/netbackup/bin/admincmd/nbstlutil</pre> Note: In NetBackup 10.3 and later, the <code>nbstlutil</code> binary is only present on NetBackup primary servers. Do not port the <code>nbstlutil</code> binary to non-primary servers.
How to use	Use <code>nbstlutil</code> to perform the following administrative actions: ■ List the status of SLP-managed images. You can print the status of SLP-processed images. Support may request this information to troubleshoot an SLP problem. ■ Cancel pending duplication operations on the selected images or image copies. When a duplication is canceled, NetBackup considers the image or image copy to be SLP complete. It does not attempt to create any more copies of the backup image. ■ Deactivate (suspend) pending and future SLP operations on selected images or image copies. NetBackup retains the image information so that processing can be resumed by the administrator at a later time. ■ Activate (resume) suspended SLP operations on selected images or image copies. See the NetBackup Commands Reference Guide for a description of all the options available for <code>nbstlutil</code>.

Table 14-1 nbstlutil details (continued)

nbstlutil information	Details
When to use	NetBackup starts a duplication session every five minutes to copy data from a backup operation for a duplication operation. Five minutes is the default frequency of the Image processing interval parameter in the SLP Parameters host properties. If the copy fails, the next four duplication sessions retry the copy. If the copy fails all four times, the copy is retried every two hours until it succeeds. Two hours is the default frequency of the Extended image retry interval parameter in the SLP Parameters host properties. Use the <code>nbstlutil</code> command in the case of a hardware problem that may require more than 15 minutes to resolve. That is, the problem may take longer to resolve than four duplication sessions five minutes apart. For example, a duplication job fails because the library has a hard failure. It may take longer than two hours to repair the library. The administrator may not want duplication jobs to begin every two hours. Use the <code>nbstlutil</code> command to inactivate the SLP while the library is repaired. When ready, the SLP can be activated and duplication jobs can begin. Note: Once the job is reactivated, the administrator may want to temporarily change the Extended image retry interval parameter to one hour to begin duplication jobs sooner.

Storage operations

This chapter includes the following topics:

- [Operation types in a storage lifecycle policy](#)
- [Backup operation in an SLP](#)
- [Backup From Snapshot operation in an SLP](#)
- [Duplication operation in an SLP](#)
- [Import operation in an SLP](#)
- [Index From Snapshot operation in an SLP](#)
- [Replication operation in an SLP](#)
- [Snapshot operation in an SLP](#)
- [Creating a hierarchy of storage operations in a storage lifecycle policy](#)

Operation types in a storage lifecycle policy

The **Operation** selections are the instructions in the storage lifecycle policy. The following topics describe the purpose of each operation.

- See [“Backup From Snapshot operation in an SLP”](#) on page 609.
- See [“Duplication operation in an SLP”](#) on page 610.
- See [“Index From Snapshot operation in an SLP”](#) on page 612.
- See [“Replication operation in an SLP”](#) on page 615.
- See [“Snapshot operation in an SLP”](#) on page 616.

Backup operation in an SLP


Use the **Backup** operation in a storage lifecycle policy to create a backup. All **Backup** operations in a single storage lifecycle policy must be on the same media server.

A **Backup** operation creates a tar-formatted image. To create a snapshot image, select a **Snapshot** operation.

Table 15-1 Backup operation characteristics

Characteristic	Description
Destination storage	The selection must be a backup storage unit or a backup storage unit group (see the following note). The selection cannot be a snapshot storage unit or a snapshot storage unit group. Note: If creating an SLP for Auto Image Replication, do not select a storage unit group. Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group. See “About NetBackup Auto Image Replication” on page 961.
Child of	A Backup operation cannot serve as the child of any other operation.
Source for	A Backup operation can be the source of a Duplication operation. (See Figure 15-1 .)
Hierarchy notes	If a Backup operation appears in an SLP, it must be the first operation. An SLP can contain a maximum of four Backup operations.
Job type	A Backup operation generates a Backup job in the Activity monitor .
Window	Backup operations do not offer the option to create an SLP window. See “Window tab” on page 634.

Figure 15-1 SLP that contains a Backup operation



Backup From Snapshot operation in an SLP

Use the **Backup From Snapshot** operation to create a tar-formatted copy of the snapshot. The new copy is a backup copy. The process is sometimes referred to as a *snapdupe* job.


Table 15-2 Backup From Snapshot operation characteristics

Characteristic	Description
Destination storage	The selection must be a backup storage unit or a backup storage unit group. The selection cannot be a snapshot storage unit or a snapshot storage unit group.
Child of	A Backup From Snapshot operation must use a Snapshot operation as its source.
Source for	A Backup From Snapshot operation can be the source for a Duplication operation. (See Figure 15-2 .)
Hierarchy notes	An SLP may contain more than one Backup From Snapshot operation. If the first Backup From Snapshot operation fails with an unrecoverable error, NetBackup does not attempt the second one. Note: The SLP may contain only one Backup From Snapshot operation if the SLP is used by an NDMP policy (or a Standard or MS-Windows policy with NDMP Data Mover enabled).

Table 15-2 Backup From Snapshot operation characteristics (continued)

Characteristic	Description
Job type	A Backup From Snapshot operation generates a Backup job in the Activity Monitor. The Backup job that results from the Backup From Snapshot operation is under the control of the SLP Manager. If an SLP window is configured, the Backup job runs during the configured SLP window. If no SLP window is configured, the Backup job can run at any time; possibly outside of the backup window as configured in the backup policy. Users may experience a slight degradation in performance on the client or the client storage device while NetBackup accesses the snapshot.
Window	An SLP window can be created for a Backup From Snapshot operation. See “Window tab” on page 634.

Figure 15-2 SLP that contains a Backup From Snapshot operation




Duplication operation in an SLP

Use the **Duplication** operation to create a copy of a **Backup**, a **Backup from Snapshot**, or another **Duplication** operation. A media server performs the operation and writes the copy.

Note: Use the **Replication** operation to create a copy of a **Snapshot** operation. See [“Replication operation in an SLP”](#) on page 615.

Table 15-3 Duplication operation characteristics

Characteristic	Description
Destination storage	The selection must be a backup storage unit or a backup storage unit group. The selection cannot be a snapshot storage unit or a snapshot storage unit group.
Child of	A Duplication operation can be the child of the following operations: ■ Backup operation ■ Backup From Snapshot operation ■ A Duplication operation
Source for	A Duplication operation can be the source for a Duplication operation. (See Figure 15-3 .)
Job type	A Duplication operation generates a Duplication job in the Activity Monitor .
Window	An SLP window can be created for a Duplication operation. See “Window tab” on page 634.

Figure 15-3 SLP that contains one Backup operation and two Duplication operations

Import operation in an SLP

Use the **Import** operation as part of Auto Image Replication. An **Import** operation in an SLP indicates that the SLP is to automatically import images into the target primary domain. An SLP that contains an **Import** operation is referred to as an **Import SLP**.

Table 15-4 Import operation characteristics

Characteristic	Description
Destination storage	An Import operation can import only from a backup storage unit or a backup storage unit group. It cannot import from a snapshot storage unit or a snapshot storage unit group.
Child of	An Import operation cannot serve as the child of any other operation.
Source for	An Import operation can be the source of a Duplication operation. At least one operation in the SLP must use the Target retention retention type.
Hierarchy notes	If an SLP contains an Import operation, it must be the first in the operations list and the only Import operation.
Job type	An Import operation generates an Import job in the Activity Monitor .
Window	An SLP window can be created for an Import operation. See "Window tab" on page 634.

The **Override job priority** option can be selected. It allows administrators to specify a job priority for any import jobs which use this SLP.

Index From Snapshot operation in an SLP

The **Index From Snapshot** operation indexes the contents of existing snapshots. When NetBackup indexes a snapshot, it creates an image catalog file (.*if*) in the NetBackup catalog for each snapshot. The presence of this file assists the user when a file needs to be restored from the snapshot, as described in [Table 15-5](#).

The **Backup From Snapshot** operation also creates an image catalog file. An **Index from Snapshot** may not be required if a **Backup From Snapshot** occurs frequently enough for the restore needs in your environment. For example, if the **Backup From Snapshot** runs once per week but file restores are required daily, consider using the **Index from Snapshot**.

The actual restore of the snapshot requires that the snapshot be mounted, regardless of whether an **Index from Snapshot** has been performed or not.

Table 15-5 Restore operations

Type of restore	Description	Requirements
Live browse restore	The user navigates the directory structure to locate and select the files for restore.	No .<i>if</i> file needs to be present in the NetBackup catalog. During a live browse restore, NetBackup automatically mounts the snapshot so that the user can see what files it contains. Mounting and unmounting the snapshot can be time-consuming.

Table 15-6 Index From Snapshot operation characteristics

Characteristic	Description
Destination storage	The Index From Snapshot operation does not write data to a storage unit. However, the destination storage is needed to select the media server that to be used to access the snapshot. As a best practice, use the storage unit from the Snapshot or Replication operation that is the source for this operation.
Child of	An Index From Snapshot is a child of a Snapshot or Replication operation.
Source for	While an Index From Snapshot operation cannot be the source for any operation, a Replication operation can follow it.
Hierarchy notes	The Index From Snapshot operation can consume system resources and requires that each snapshot be mounted to create the .<i>if</i> file. See “Determining where and when the Index From Snapshot operation occurs” on page 614.
Job type	An Index From Snapshot operation generates an Index From Snapshot job in the Activity monitor .
Window	An SLP window can be created for an Index From Snapshot operation.

Consider the following items before using the **Index From Snapshot** operation:

- The **Index From Snapshot** operation is supported only in a Replication Director configuration.
- **Standard**, **MS-Windows**, **NDMP**, and **VMware** backup policy types support the use of storage lifecycle policies that contain the **Index From Snapshot** operation. However, a **Standard** or **MS-Windows** policy with **NDMP Data mover** enabled is not supported.
- The **Index From Snapshot** operation can run from a full or an incremental schedule. The file entries that are added to the .*if* file for either schedule are the full set of files since all files can be restored from that snapshot. To do so

allows for the most efficient restore, however, more space is consumed in the NetBackup catalog by the . ϵ file.

Determining where and when the Index From Snapshot operation occurs

Including the **Index From Snapshot** operation requires some consideration as the operation can consume system resources and require additional time to perform. For example, to perform the operation can require that a snapshot be mounted or that NetBackup gather content details from the file system to populate the catalog.

To help mitigate the extra resource and time that the operation may take, the system administrator can control when and where the **Index From Snapshot** operation runs:

- Use the storage lifecycle policy **Window** tab to schedule when the **Index From Snapshot** operation can run. Schedule the operation to run when it is least likely to interfere with other jobs.
See [“Window tab”](#) on page 634.
- Use the following points to determine where to position the **Index From Snapshot** operation in the SLP operations list:
 - Each NetBackup environment needs to determine where the operation works best in a specific SLP. To place the **Index From Snapshot** operation too early (toward the top of the operations list), may consume time when the restore capabilities are not needed. To place the operation toward the end of the operations list may cause the administrator to delay a restore until earlier snapshots or replications complete.
 - Use the **Index From Snapshot** operation in an SLP only once. A restore can be performed from any snapshot after one image . ϵ file is created.
 - Any operations list that includes a **Backup From Snapshot** operation does not need an **Index From Snapshot** operation. The **Backup From Snapshot** operation creates an image . ϵ file. The only exception is if the index is needed for restores before the **Backup From Snapshot** operation occurs.
 - An **Index From Snapshot** operation cannot have any dependents. An SLP cannot validate an **Index From Snapshot** operation with children.
See [Table 15-7](#) and [Table 15-8](#) for valid SLP configurations.
A **Replication** operation follows the **Index From Snapshot** operation. But it is a child of the **Snapshot** operation, not a child of the **Index From Snapshot** operation.
To add a **Replication** operation after an **Index From Snapshot** operation, select the **Snapshot** operation, then click **Add**.

Table 15-7 Example 1 of a valid placement of the Index From Snapshot operation

Snapshot	STU_Primary1
Replication	STU_2
Index from Snapshot	Any available

Table 15-8 Example 2 of a valid placement of the Index From Snapshot operation

Snapshot	STU_Primary1
Index from Snapshot	Any available
Replication	STU_2

Replication operation in an SLP

Use the **Replication** operation for the following types of replication:

- NetBackup Replication Director to replicate a snapshot.
See [“About NetBackup Replication Director”](#) on page 1002.
- NetBackup Auto Image Replication to replicate a backup.
See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 977.

Table 15-9 Replication operation characteristics

Characteristic	Description
Replication target	Under Destination storage attributes: ■ For Replication Director, select the Storage that is configured to contain replicated snapshots. ■ For Auto Image Replication, choose to either: ■ Replicate the backup to storage servers in all target NetBackup domains. ■ Replicate the backup to a specific master server in a specific domain. This type of Auto Image Replication is known as targeted A.I.R.

Table 15-9 Replication operation characteristics (*continued*)

Characteristic	Description
Child of	A Replication operation can be the child of any of the following operations: ■ Snapshot operation for NetBackup Replication Director to replicate a snapshot. ■ Another Replication operation. ■ Backup operation for NetBackup Auto Image Replication.
Source for	A Replication operation can be the source for the following operations: ■ Replication ■ Backup From Snapshot See “Backup From Snapshot operation in an SLP” on page 609.
Job type	A Replication operation generates a Replication job in the Activity Monitor .
Window	An SLP window can be created for a Replication operation. See “Window tab” on page 634.

Snapshot operation in an SLP

A **Snapshot** operation creates a point-in-time, read-only, disk-based copy of data. NetBackup provides several types of snapshots, depending on the device where the snapshot occurs.

Use a **Snapshot** operation as the first operation in a storage lifecycle policy for a NetBackup Replication Director configuration.

Table 15-10 Snapshot operation characteristics


Characteristic	Description
Destination storage	The following topics describe the types of snapshot storage units that can be used as the storage for a snapshot operation: ■ See “Primary snapshot storage unit” on page 618. ■ See “Primary + Replication source snapshot storage unit” on page 618. ■ See “Replication source + Replication target snapshot storage unit” on page 619. ■ See “Replication target snapshot storage unit” on page 619. ■ See “Replication source + Replication target + Mirror snapshot storage unit” on page 620. A Storage unit selection is necessary in the following situations: ■ If the Snapshot is to be used by a subsequent Replication operation. The storage unit that is specified for the Snapshot operation must be a snapshot-capable storage unit that represents the primary storage. ■ If the SLP contains only one operation and that is a Snapshot operation, specify a storage unit. NetBackup uses that storage unit to determine which media server to use to launch the snapshot job. If neither situation applies to the SLP, the administrator can select No storage unit. NetBackup uses the storage unit that is selected for the Backup From Snapshot operation.
Child of	A Snapshot operation cannot serve as the child of any other operation.
Source for	A Snapshot operation can be the source for the following operations: ■ Backup From Snapshot ■ Index From Snapshot ■ Replication operation
Hierarchy notes	If a Snapshot operation appears in an SLP, it must be first in the operations list.
Job type	A Snapshot operation generates a Snapshot job in the Activity Monitor .
Window	Snapshot operations do not offer the option to create an SLP window. See “Window tab” on page 634.

Primary snapshot storage unit

A snapshot operation can use a **Primary** snapshot storage unit. That is, the storage unit represents a disk pool that contains the volumes that have only the **Primary** property set.

Figure 15-4 shows an SLP that contains one primary-only **Snapshot** operation, one **Backup From Snapshot** operation, and one **Duplication** operation. The **Backup From Snapshot** operation is used to create a backup from the snapshot on the primary-only **Snapshot** operation. After the backup is created, it is duplicated to a **Duplication** operation.

Figure 15-4 SLP that contains a Snapshot operation, a Backup From Snapshot operation, and a Duplication operation




Primary + Replication source snapshot storage unit

An SLP operation can use a **Primary + Replication source** snapshot storage unit. That is, the storage unit represents a disk pool that contains volumes that have both the **Primary** property and the **Replication source** property set.

Figure 15-5 shows an SLP that contains a **Primary + Replication source** snapshot storage unit as one operation and one **Replication target** snapshot storage unit as another operation. The **Primary + Replication source** storage unit can replicate to the **Replication target** storage unit.

Figure 15-5 SLP that contains a Snapshot operation and a Replication operation




Replication source + Replication target snapshot storage unit

An SLP operation can use a snapshot storage unit that represents a disk pool that contains volumes that have the following properties: **Replication source** and **Replication target**.

A snapshot storage unit with these properties can serve as both the **Replication source** for another operation in the SLP, and as the **Replication target** for another operation in the SLP.

Figure 15-6 SLP that contains a Snapshot operation and two Replication operations




Replication target snapshot storage unit

An SLP operation can use a snapshot storage unit that represents a disk pool that contains volumes that have only the **Replication target** property set.

An operation with this property can serve only as a **Replication target** for another operation in the SLP. It cannot serve as source for a replica, but it can serve as the source for a **Duplication** operation.

Figure 15-7 SLP that contains a Snapshot operation, a Replication operation, a Backup From Snapshot operation, and a Duplication operation



Replication source + Replication target + Mirror snapshot storage unit

An SLP can use a snapshot storage unit that represents a disk pool that contains volumes that have the following properties: **Replication source**, **Replication target**, and **Mirror**.

An operation with these properties can serve as both:

- A **Replication source** in a cascading configuration.
- A mirrored **Replication target** in a cascading configuration. A mirrored **Replication target** must have a forced **Mirror** retention type.

Replication target + Mirror snapshot storage unit

An SLP can use a snapshot storage unit that represented a disk pool that contains volumes that have the following properties: **Replication target** and **Mirror**.

A mirrored **Replication target** must have a forced **Mirror** retention type.

Creating a hierarchy of storage operations in a storage lifecycle policy


The list of operations in the storage lifecycle policy acts as a list of instructions to NetBackup about the data that the SLP protects. In some cases, one operation

depends on another operation. For example, a snapshot may serve as the source for a replication. Or, a backup may serve as the source of a duplication.

The operation hierarchy in the SLP represents a parent and child relationship.

One copy can be the source for many other copies. [Figure 15-8](#) shows how after the first copy is created, all subsequent copies can be made locally from that source, without tying up network resources.

Figure 15-8 Hierarchical operations



Changing the location of an operation in the hierarchy changes the storage unit that serves as the source for the subsequent copies. Changing the hierarchy cannot change the operation type. (For example, change a backup operation into a duplication operation.)

Retention types for SLP operations

This chapter includes the following topics:

- [Retention types for storage lifecycle policy operations](#)
- [Capacity managed retention type for SLP operations](#)
- [Expire after copy retention type for SLP operations](#)
- [Fixed retention type for SLP operations](#)
- [Maximum snapshot limit retention type for SLP operations](#)
- [Mirror retention type for SLP operations](#)
- [Target retention type for SLP operations](#)

Retention types for storage lifecycle policy operations

The **Retention type** for an operation in a storage lifecycle policy determines how long the data is kept on that storage media.

[Table 16-1](#) describes which retention types are valid selections for the various operations.

Table 16-1 Operation and retention type configurations

Retention type	Backup operation	Snapshot operation	Replication operation	Backup From Snapshot operation	Duplication operation
Fixed	Valid	Valid	Valid	Valid	Valid
Expire after copy	Valid	Valid	Valid	Valid	Valid
Maximum Snapshot limit	Invalid	Valid; SLP honors the policy setting.	Invalid	Invalid	Invalid
Mirror	Invalid	Invalid	Valid for snapshot storage only	Invalid	Valid for snapshot storage only
Target retention	Invalid	Invalid	Valid if the first operation in the SLP is an Import and if the storage is of the backup type.	Invalid	Valid if the first operation in the SLP is an Import.
Capacity managed	Valid; AdvancedDisk default; set on the storage server.	Invalid	Invalid	Invalid	Valid; AdvancedDisk default; set on the storage server.

Note: Retention is not associated with the **Index From Snapshot** operation because the operation does not create any copy.

Mixing retention types

It is not recommended to allow capacity-managed images and fixed-retention images to be written to the same volume in a disk storage unit. The volume may fill with fixed-retention images and not allow the space management logic to operate as expected.

Keep in mind the following points when configuring SLP operations or selecting the storage location for a policy:

- All SLPs that write to a volume in a disk storage unit should write images of the same retention type: fixed or capacity-managed.

- Do not write images both to a volume in a disk storage unit within an SLP and to the same volume (by the storage unit) directly from a policy.
- Mark all disk storage units that are used with SLPs as **On demand only**.
- Check any storage unit groups to make sure that fixed and capacity-managed images cannot be written to the same volume in a disk storage unit.

Capacity managed retention type for SLP operations

A **Capacity managed** operation means that NetBackup automatically manages the space on the storage, based on the **High water mark** setting for each volume. **Capacity managed** is not available to tape storage units since tape capacity is considered to be infinite.

The **High water mark** and **Low water mark** settings on the disk storage unit or disk pool determine how the space is managed.

See [“High water mark storage unit setting”](#) on page 561.

See [“Low water mark storage unit setting”](#) on page 562.

An image copy with a **Capacity Managed** retention is not eligible for expiration until its dependent copies have been created.

If space is needed for new images, NetBackup removes expired backup images from a capacity-managed disk volume in two passes, as follows:

Pass one	NetBackup removes any backup images that are past the Desired cache period setting. NetBackup removes images until the low water mark is reached or all images that are past the Desired cache period are removed.
Pass two	Pass two processing is initiated if the outcome of the pass one processing is one of the following: ■ The disk pool remains over the high water mark. ■ The number of volumes in the disk pool under the high water mark is less than the number of media servers that access the disk pool.

NetBackup removes images until the low water mark is reached or all images that are not past the **Desired cache period** are removed.

An image may be deleted if it has not been duplicated for all operations in a storage lifecycle policy. If the operating system time is past the date that matches the longest retention period for an image, the image is eligible for deletion.

To see exactly when the storage reaches the low water mark value is difficult. A backup can occur at the same time as the expiration process occurs. After the backup is complete, the low water mark may be slightly greater than its lowest possible value.

The retention period for capacity managed storage is not assured as it is for a fixed retention period. The **Desired cache period** becomes a target that NetBackup tries to maintain. If the space is not required, the backup data could remain on the storage longer than the **Desired cache period** indicates.

Rules and recommendations for using the Capacity Managed retention type

Use the following recommendations and rules when configuring storage operations or when selecting the storage location for a policy:

- It is not recommended to allow **Capacity Managed** images and **Fixed** retention images to be written to the same volume in a disk storage unit. The volume may fill with fixed-retention images and not allow the space management logic to operate as expected.
- All SLPs that write to a volume in a disk storage unit should write images of the same retention type: **Fixed** or **Capacity Managed**.
- Do not write images both to a volume in a disk storage unit within a storage lifecycle policy and to the same volume (by the storage unit) directly from a policy.
- Mark all disk storage units that are used with SLPs as **On demand only**.
- Check any storage unit groups to make sure that fixed and capacity-managed images cannot be written to the same volume in a disk storage unit.

Capacity managed retention type and disk types that support SIS

Capacity managed is selectable for any disk storage unit that is allowed in an SLP. However, for the disk types that support single-instance storage (SIS), **Capacity managed** functions to various degrees. In order for **Capacity managed** to operate, NetBackup must know how much space a backup image uses. With SIS enabled on the storage unit, NetBackup cannot know exactly how much space a particular backup image occupies.

The following storage unit configurations use SIS:

- **Media Server Deduplication Pool** storage units
- Some OpenStorage storage units, depending on the vendor characteristics.

Expire after copy retention type for SLP operations

The **Expire after copy** retention indicates that after all direct (child) copies of an image are successfully duplicated to other storage, the data on this storage is expired. The last operation in the SLP cannot use the **Expire after copy** retention type because no subsequent copy is configured. Therefore, an operation with this retention type must have a child.

It is not recommended that you enable **Expire after copy** retention for any storage units that are to be used with SLPs with either of the following: Accelerator or synthetic backups. The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.

Note: Although synthetic backups do support the use of storage lifecycle policies, SLPs cannot be used for the multiple copy synthetic backups method.

If a policy is configured to use an SLP for the backup, the retention that is indicated in the SLP is the value that is used. The **Retention** attribute in the schedule is not used.

An image copy with an **Expire after copy** retention is expired as soon as all of its direct child copies have been successfully created. Any mirrored children must also be eligible for expiration.

Fixed retention type for SLP operations

The **Fixed** retention indicates that the data on the storage is retained for the specified length of time, after which the backups or snapshots are expired.

An image copy with a **Fixed** retention is eligible for expiration when all of the following criteria are met:

- The **Fixed** retention period for the copy has expired.
- All child copies have been created.
- All child copies that are mirror copies are eligible for expiration.

The **Fixed** retention period is always marked from the original backup time of the image. For example, if a tape device is down, causing a 2-day delay in creating a duplicate tape copy, the expiration time of the duplicate copy is not different due to the 2-day delay. The expiration time of the duplicate copy is still x days from the

time that the original backup was completed. It does not matter when the copy was created.

Maximum snapshot limit retention type for SLP operations

The **Maximum snapshot limit** determines the maximum number of snapshots that can be stored for a particular policy and client pair.

When the maximum is reached, the next snapshot causes the oldest job-complete snapshot to be deleted. A snapshot job is considered to be complete once all of its configured dependent copies are complete. (Dependent copies are created as a result of **Backup From Snapshot**, **Index From Snapshot**, or **Replication** operations.) The practice is referred to as *rotation*. This retention type applies only to snapshots, and not to backups.

For example, Policy P1 contains two clients: C1 and C2. After the policy runs four times, it creates four snapshot images for C1 and four images for C2. If the **Maximum snapshot limit** is set to four, when the policy runs for the fifth time, NetBackup deletes the first snapshot that was created for both C1 and C2 to accommodate the fifth snapshot.

The **Maximum Snapshots** parameter determines the maximum number of snapshots. To configure this setting select the **Snapshot options** in the backup policy.

See [“Snapshot Client and Replication Director \(policy attributes\)”](#) on page 728.

Mirror retention type for SLP operations

A mirror replica of a snapshot is eligible for expiration as soon as:

- All immediate child copies are successfully created.
- All immediate child copies that are mirrors are eligible for expiration.

The selection of the **Mirror** retention indicates that NetApp volume SnapMirror is to be used as the replication method. If any non-mirror retention type such as **Fixed** or **Expire after copy** is selected for the **Replication** operation, the NetApp SnapVault replication method is used.

In mirror replication, the replica copy is dependent on the existence of the source. (The source can be the original snapshot or another replica.) Therefore, the retention of the replica depends on the retention of the source. If the source is deleted, the mirror is automatically deleted.

In non-mirror replication, the replica is independent of the source and can have an independent retention. If the source is deleted, the non-mirror replica is not affected and can be used longer than the source. Or, if the replica is deleted first, it is not recreated and the source can be kept longer than the replica.

Target retention type for SLP operations

This setting is used in Auto Image Replication in an Import storage lifecycle policy. Every Import SLP must have at least one operation with a **Target retention**.

The **Target retention** is enforced at the target domain, but the actual retention for the data is specified by the administrator at the source domain.

Target retention indicates that the data at the target master shall use the expiration date that was imported with the image. The date is fixed because the copy must have a fixed retention.

Similar to the **Fixed** retention, an image copy with a **Target retention** retention is eligible for expiration when all of the following criteria are met:

- The **Fixed** retention period for the copy has expired.
- All child copies have been created.
- All child copies that are mirror copies are eligible for expiration.

See [“About NetBackup Auto Image Replication”](#) on page 961.

Storage lifecycle policy options

This chapter includes the following topics:

- [Storage lifecycle policy settings](#)
- [SLP settings](#)
- [Storage lifecycle policy Validation report tab](#)

Storage lifecycle policy settings

The following table describes the settings for storage lifecycle policies.

Figure 17-1 Storage lifecycle policy tab

Storage Lifecycle Policy

Storage lifecycle policy Validation report

Storage lifecycle policy name: SLP_1_snapshot

Data classification: No data classification

Priority for secondary operations: 0

A higher number is greater priority.

+ Add

Operation	Storage	Storage type	Volume pool	Media owner	Retention type	Retention period
☐ Snapshot	No Storage Unit				Maximum Snapshot Limit	2 weeks
☐ Backup From Snapshot	stu_adv	AdvancedDisk			Fixed	2 weeks

2 Records

State of secondary operation processing

To find impact on policies associated with this SLP due to change in configuration click here.

Cancel Create

Table 17-1 Storage lifecycle policy tab

Setting	Description
Storage lifecycle policy name	The Storage lifecycle policy name describes the SLP. The name cannot be modified after the SLP is created.
Data classification	The Data classification defines the level or classification of data that the SLP is allowed to process. The dropdown menu contains all of the defined classifications as well as the Any classification, which is unique to SLPs. The Any selection indicates to the SLP that it should preserve all images that are submitted, regardless of their data classification. This setting is available for SLP configuration only and is not available to configure a backup policy. In an Auto Image Replication configuration where the master server domains run different versions of NetBackup, see the following topic for special considerations: See “About the storage lifecycle policies required for Auto Image Replication” on page 977. The Data classification is an optional setting. One data classification can be assigned to each SLP and applies to all operations in the SLP. If a data classification is selected (other than Any), the SLP stores only those images from the policies that are set up for that data classification. If no data classification is indicated, the SLP accepts images of any classification or no classification. The Data classification setting allows the NetBackup administrator to classify data based on relative importance. A classification represents a set of backup requirements. When data must meet different backup requirements, consider assigning different classifications. For example, email backup data can be assigned to the silver data classification and financial data backup may be assigned to the platinum classification. A backup policy associates backup data with a data classification. Policy data can be stored only in an SLP with the same data classification. Once data is backed up in an SLP, the data is managed according to the SLP configuration. The SLP defines what happens to the data from the initial backup until the last copy of the image has expired.
Priority for secondary operations	The Priority for secondary operations option is the priority that jobs from secondary operations have in relationship to all other jobs. The priority applies to the jobs that result from all operations except for Backup and Snapshot operations. Range: 0 (default) to 99999 (highest priority). For example, you may want to set the Priority for secondary operations for a policy with a gold data classification higher than for a policy with a silver data classification. The priority of the backup job is set in the backup policy on the Attributes tab.

Table 17-1 Storage lifecycle policy tab (*continued*)

Setting	Description
Operations	Use the Add child, Edit, and Delete options to create a list of operations in the SLP. An SLP must contain one or more operations. Multiple operations imply that multiple copies are created. The list also contains the columns that display information about each operation. Not all columns are displayed by default.
Move	Use the move option to change the hierarchy of the source for each copy. One copy can be the source for many other copies.
State of Secondary Operation Processing	The Active and Postponed options refer to the processing of all duplication operations in the SLP. Note: These options apply to any duplication operations that create tar-formatted images. For example, those created with <code>bpduplicate</code>. These options do not affect the images that are duplicated as a result of OpenStorage optimized duplication or NDMP. Nor do they affect if one or more destination storage units are specified as part of a storage unit group. ■ Enable Active to let secondary operations continue as soon as possible. When an operation is changed from Postponed to Active, NetBackup continues to process the images, picking up where it left off when secondary operations were made inactive. ■ Enable Postponed to postpone the secondary operations for the entire SLP. Postponed does not postpone the creation of duplication jobs, it postpones the creation of images instead. The duplication jobs continue to be created, but they are not run until secondary operations are active again. All secondary operations in the SLP are inactive indefinitely unless the administrator selects Active or until the Until option is selected and an activation date is indicated.
Check for conflicts with backup policies option	Select this option to see how changes to this SLP can affect the policies that are associated with this SLP. The option generates a report that displays on the Validation report tab. This option performs the same validation as the <code>-conflict</code> option performs when used with the <code>nbstl</code> command.

SLP settings

When you create or edit a storage operation, the following tabs are available:

- **Properties** tab

The following topic describes the options in the **Properties** tab:
See [“Properties tab for an SLP operation”](#) on page 632.

- **Window** tab

Create a window during which a secondary operation can run. The following topic describes the options in the **Window** tab:

See [“Window tab”](#) on page 634.

Properties tab for an SLP operation

[Table 17-2](#) describes the options available to configure for the various operations in a storage lifecycle policy. Not all options are available for each operation.

Table 17-2 Properties tab for an SLP operation

Setting	Description
Source storage	Indicates the storage unit that is the source for the operation. The Source storage displays for the following operation types: Backup From Snapshot, Replication, Duplication, and Index From Snapshot.
Operation	The Operation selection determines which options display. See “Operation types in a storage lifecycle policy” on page 607. See “About writing multiple copies using a storage lifecycle policy” on page 639.
Destination storage	Select the destination storage that you want to use. Storage unit groups may contain no BasicDisk or disk staging storage units. A storage lifecycle policy can point to a storage unit group that contains a BasicDisk storage unit. However, NetBackup does not select BasicDisk storage units from a storage group for a lifecycle policy. Storage units or storage unit groups may appear in more than one lifecycle. Storage units or storage unit groups may be used in a storage lifecycle while also being used as standalone units.
Destination storage attributes	Indicates that the copy of the image is to be created in a different primary server domain. All replication target storage servers Sends the backups to all the replication target storage servers. Target primary server The target primary server manages the storage where the image is to be copied. When this option is selected for a Replication operation, the operation becomes an operation for Auto Image Replication.
Volume pool	The Volume pool option is enabled for tape storage units or virtual tape libraries (VTL).
Media owner	A Media owner is a group of NetBackup servers that are used for a common purpose. By specifying a Media owner, you allow only those media servers to write to the media on which backup images for a policy are written.

Table 17-2 Properties tab for an SLP operation (*continued*)

Setting	Description
Retention type	Select a Retention type from the following options: ■ See “Capacity managed retention type for SLP operations” on page 624. ■ See “Expire after copy retention type for SLP operations” on page 626. ■ See “Fixed retention type for SLP operations” on page 626. ■ See “Maximum snapshot limit retention type for SLP operations” on page 627. ■ See “Mirror retention type for SLP operations” on page 627. ■ See “Target retention type for SLP operations” on page 628. See “Retention types for storage lifecycle policy operations” on page 622.
Retention period	Specifies how long NetBackup retains the backup or copy. To set the retention period, select a time period (or level) from the list. When the retention period expires, NetBackup deletes information about the expired backup or copy. After expiration, the files are unavailable for restores. For example, if the retention is set to 2 weeks, data can be restored from a backup that this schedule performs for 2 weeks after the backup. See “Retention periods properties” on page 152.
Duplication	Alternate read server This option specifies the name of the server that is allowed to read a backup image originally written by a different media server. This option is available for Duplication operations only.
Preserve multiplexing	The Preserve multiplexing option is available for the duplication operations that use tape media or virtual tape libraries (VTL). If the backup to be duplicated is multiplexed and you want the backups to remain multiplexed, select this option. To preserve multiplexing significantly improves the performance of duplication jobs because it eliminates the need to request the write-side duplication media for every image.
Override job priority	The Override job priority option is available for an Import operation. The job priority that is indicated is the job priority for any import jobs which use this storage lifecycle policy.
Postpone creation of this copy until the source copy is about to expire	Enable this option to defer the job until the source for the duplication is about to expire. When this option is enabled, the job begins 4 hours before the source is to expire. This default can be changed by changing the Deferred duplication offset time in the SLP Parameters host properties. See “SLP settings properties” on page 168.

Table 17-2 Properties tab for an SLP operation (*continued*)

Setting	Description
Advanced settings	Windows close behavior If a window closes and the jobs for an SLP have not completed, NetBackup attempts to suspend the images that are in progress. When the window reopens, NetBackup resumes those jobs at the point they were when suspended. Not all images can be suspended. The image must be the result of a duplication job where both the source and target of the duplication resides on either an AdvancedDisk or Media Manager storage unit. In addition, the duplication job must meet the following conditions: ■ The storage units must not be part of a storage unit group. ■ The duplications were not created using optimized duplication, NDMP duplication, or OpenStorage duplication. See "Duplication operation in an SLP" on page 610. Note: The closing of the window does not stop preliminary operations for some jobs. For example, NetBackup continues to extend the catalog for Exchange Granular Recovery after the window closes, but does not start duplications. Images that result from all other operations (such as Replication operations), are not suspended. Select the Advanced dropdown list to display the Windows close behavior options. The selections apply to the images that NetBackup does not suspend automatically. Select what NetBackup should do if images do not complete by the time the window closes and if the images cannot be suspended: ■ Finish processing the active images The window closes, but NetBackup continues to process the active images until they are finished. NetBackup does not begin to process any other images until the window reopens. ■ Cancel the processing of the active images The window closes and NetBackup immediately stops processing the active images. When the window reopens, NetBackup begins to process the images where it left off.

Window tab

The **Window** tab appears for secondary operations in a storage lifecycle policy.

Creating a window for a secondary operation is optional. However, creating a window can better define when the job for the operation can run. In this way, the job from a secondary operation does not interfere with jobs of a higher priority, such as backup jobs. Without a window defined, the job for an operation can run at any time, on any day.

Table 17-3 Window tab

Setting	Description
Windows options drop down list:	You can either assign an existing window to the operation or create a new window for the operation. To use an existing window, select Select from saved windows and then select a window from the list of window names. To create a new window, select Create new for this operation to use. Then enter a name for the new window.
Impact report	This panel lists the names of the storage lifecycle policies that currently use the window. The Impact report also lists the operation that uses the window, and the source and the destination storage for the operation.
Start window tab	The Start window grid is grayed out and cannot be modified if the Default_24x7_Window is selected. The Start window grid is active if a saved window is selected or when a new window is created. If the Start window grid is changed for a saved window, click the Impact report option to display information about other operations in other SLPs that use the window. See “Creating a new window for a storage lifecycle policy operation” on page 635.
Exclude datestab	Use the Exclude dates tab to exclude specific dates from a window. See “Excluding days from a window for a storage lifecycle policy operation” on page 636.

Creating a new window for a storage lifecycle policy operation

To create a new window for an SLP operation

- 1 In the **Window** tab of the storage operation dialog box, enable **Create new**.
- 2 Select the **Start Window** tab.
- 3 The days of the week appear along the left side of the grid. The time of day appears along the top of the grid in 24-hour time.

To change the increments available for selecting start times or end times, change the value in the **Resolution** field.
- 4 Indicate the opening and closing times of the window for each day. The following lists several methods to do so:

- Drag the cursor along the **Start window** grid on each day you want the window to open and close.
- Use the dropdown menus to select a **Start day** and an **End day**. Then select a **Start time** and an **End time**.
- Use the dropdown menu to select a **Start day** and the **Duration days** of the window. Adjust the **Start time** for your environment.

To create multiple time windows:

- First, create one window.
- Select **Duplicate**.
The window is duplicated to any days without existing schedules. Duplication stops when it reaches a day that already contains a window.
- On days that you do not want the time window to be open, select the window and select **Delete**.

5 You can perform the following actions:

To change the start time or end time Adjust the **Start time** or **End time**.

To delete a time window Select a time window and click **Delete**.

To delete all the time windows Click **Clear**.

To erase the last action Click **Undo**.

6 Click **OK** to save the window and the operation.

Excluding days from a window for a storage lifecycle policy operation

Use the **Exclude dates** tab to exclude specific days from a window. If a day is excluded from a window, jobs do not run on that day. The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

To exclude a day from the storage lifecycle policy window

- 1** In the **Window** tab, select the name of an existing window from the drop-down menu.
- 2** Select the **Exclude dates** tab.
- 3** Use one or more methods to indicate the days to exclude:
 - Select the days on the 3-month calendar that you want to exclude. Use the drop-down lists at the top of the calendar to change the months or year.

- To indicate **Recurring week days**:
 - Select **Select all** to select all of the days in every month for every year.
 - Select **Deselect all** to remove all existing selections.
 - Select a box in the matrix to select a specific day to exclude for every month.
 - Select the column head of a day of the week to exclude that day every month.
 - Select the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to exclude that week every month.
 - To indicate **Recurring days of the month**:
 - Select **Select all** to select all of the days in every month.
 - Select **Deselect all** to remove all existing selections.
 - Select a box in the matrix to select that day to exclude each month.
 - Select **Last day** to exclude the last day of every month.
 - To indicate **Specific dates**:
 - Select **Add**. Enter the month, day, and year in the **Date selection**. The date appears in the **Specific dates** list.
 - To delete a date, select the date in the list. Select **Delete**.
- 4 Add additional dates as necessary, and then select **Save**.

Storage lifecycle policy Validation report tab

NetBackup validates the SLP when it is first created and whenever it is changed. The **Validation report** tab lists the conflicts between the proposed changes to the operations in a storage lifecycle policy and any backup policy that uses the SLP.

Likewise, when a policy is created that indicates an SLP as the **Policy storage**, a similar validation report may display. The report lists any conflicts between the policy and the SLP that it has indicated.

The conflicts that are listed must be resolved in order for a job that references the SLP to run successfully. Use the **Troubleshooter** or the online Help on this tab for a list of common status codes that result from SLP conflicts.

Note: The `Request has timed out` message may appear in environments with very busy servers.

To increase the timeout to account for the latency in connectivity, increase the `NBJAVA_CORBA_DEFAULT_TIMEOUT` value in the following files:

On Windows: The `setconf.bat` configuration file
(`Install_path\java\setconf.bat`).

On UNIX: The `nbg.conf` configuration file (`/usr/opensv/java/nbg.conf`).

The report checks for the following conflicts between the selected SLP and the backup policies that use it:

- The data classification in the storage lifecycle policy does not match that in the referencing backup policies.
- The SLP contains a **Snapshot** operation, but the referencing backup policies do not have **Perform snapshot backups** enabled.
- The SLP does not contain a **Snapshot** operation, but the referencing backup policies have the **Perform snapshot backups** enabled.
- The SLP contains a **Snapshot** operation, but the referencing backup policies cannot enable the **Retain snapshots for Instant Recovery** option.
- The policy is of the **NBU-Catalog** backup type, but the SLP configuration does not indicate a **CatalogBackup** volume pool.
- The policy is not of the **NBU-Catalog** backup type, but the SLP configuration indicates a **CatalogBackup** volume pool

Using a storage lifecycle policy to create multiple copies

This chapter includes the following topics:

- [About writing multiple copies using a storage lifecycle policy](#)
- [How the order of the operations determines the copy order](#)
- [About ensuring successful copies using lifecycles](#)

About writing multiple copies using a storage lifecycle policy

A storage lifecycle policy can be used to create multiple copies of backups and snapshots.

NetBackup permits only one method to create multiple copies to be in use at one time. Use only one of the following methods:

- Enable the **Multiple copies** option in a policy configuration.
If a policy has the **Multiple copies** option enabled, the policy cannot select a storage lifecycle policy as the **Policy storage**.
See [“Multiple copies \(schedule attribute\)”](#) on page 747.
- Add multiple **Backup** operations or one or more **Duplication** or **Replication** operations to the operations list of the SLP.
See [“SLP settings”](#) on page 631.

The same criteria for creating copies applies to both methods.

The following topics are considerations when storage lifecycle policies are used to create multiple copies.

How the order of the operations determines the copy order

The order in which the operations appear in a storage lifecycle policy determines the copy number.

For example, the following lifecycle is configured to create three copies:

- Two copies as part of two different backup operations.
- One copy as part of a duplication operation.

To make sure that copy 1 is written to disk, place the **Backup** operation that writes to a disk storage unit before the **Backup** operation that writes to a tape storage unit.

Figure 18-1 Operation order determines copy order

	+ Add	
	☐	Operation Storage
Copy 1 on disk —	☐	Backup Disk_1
Copy 2 on tape —	☐	Duplication Tape_1
Copy 3 on disk —	☐	Backup Disk_2

About ensuring successful copies using lifecycles

The process to create copies as part of a storage lifecycle policy differs from the process to create copies as set up in a policy. The policy's **Configure multiple copies** settings include the option to **Fail all copies**. That option means that if one copy fails, the remaining copies can be set to either continue or fail.

In an SLP, all copies must be completed. An SLP initially tries three times to create a copy. If no copy is created, NetBackup continues to try, but less frequently.

The successful completion of copies is important because an SLP does not allow a copy to be expired before all copy operations in the SLP are complete. NetBackup changes the retention period of a copy to Infinity until all copies are created. After all copies are complete, the retention returns to the level as set in the policy.

To create successful copies, a **Backup** operation may be required to duplicate a backup onto the storage unit for another **Backup** operation.

Consider the following example: The operations list for an SLP contains two **Backup** operations to two storage units (BU_1, BU_2) and three **Duplication** operations.

The backup to BU_1 is successful, but the backup to BU_2 is unsuccessful.

To fulfill the backup on BU_2, NetBackup creates a duplication job from BU_1 to BU_2. The duplication job is in addition to the jobs that are run for the three duplication operations.

Duplication jobs can be controlled by using the `nbstlutil` command.

See [“Lifecycle operation administration using the nbstlutil command”](#) on page 605.

Storage lifecycle policy versions

This chapter includes the following topics:

- [About storage lifecycle policy versions](#)
- [Storage lifecycle changes and versioning](#)
- [When changes to storage lifecycle policies become effective](#)
- [Deleting old storage lifecycle policy versions](#)

About storage lifecycle policy versions

Once a storage lifecycle policy is configured, it runs according to a single configuration or definition. The definition affects the operations once they begin to run as well as the copies once the image is in process.

The ability to create SLP versions lets administrators safely modify a definition without waiting until all of the copies that are associated with the SLP have been processed. Each copy that an SLP manages is tagged with the SLP name and the SLP version number. These two attributes are written into the image header, in the NetBackup image catalog. Whenever an administrator creates or changes an SLP, NetBackup creates a new version (between 0 and n). New jobs use the most recent SLP version.

When a new job is submitted to the Activity Monitor, the job is tagged with the most recent SLP version number. The processing of a copy that is associated with a version remains fixed according to that version of the SLP definition. It is fixed at job time and does not change, unless the administrator uses the `nbstl` command to modify an existing version.

An SLP version remains as long as there are any incomplete images that refer to the version.

Storage lifecycle changes and versioning

Administrators can make changes to a storage lifecycle policy in one of the following ways:

- Using the NetBackup web UI.
Any change that an administrator makes to an SLP using the NetBackup web UI creates a new SLP version. The new version is created when the changes to the SLP are committed or saved. The NetBackup web UI always displays the most recent version.
- Using the `nbstl` command.
If an administrator uses `nbstl` to change an SLP, `nbstl` creates a new version by default.
However, the `nbstl` command contains options to view different versions and to modify the definitions of existing SLP versions without creating a new version. The options are as follows:

<code>-all_versions</code>	Use to display all versions of an SLP definition. Without specifying this option, only the most recent version is displayed by default.
<code>-version <i>number</i></code>	Use to display a specific version.
<code>-modify_current</code>	Use with most <code>nbstl</code> configuration options to make changes to the current SLP version without creating a new version. Knowing the current version number is not necessary if this option is used.
<code>-modify_version</code> <code>-version <i>number</i></code>	Use with most <code>nbstl</code> configuration options to make changes to a specific version without creating a new version.

Use `-modify_current` or `-modify_version` to change any of the following configuration options:

<code>-dp</code>	The duplication priority.
<code>-residence</code>	The storage unit to be used for each operation.
<code>-pool</code>	The volume pool for each operation.
<code>-server_group</code>	The server group for each operation.

<code>-rl</code>	The retention level for each operation.
<code>-as</code>	The alternate read server for each operation.
<code>-mpx</code>	The preserve multiplexing option for duplication copies.

Some fields require values for all of the operations in the SLP. Make sure that the number of values that are specified for the fields matches the existing operation count.

For example, in an SLP that contains three operations, to change the value of one, a value must be given for all three operations. Note that the values for all three operations are replaced. To change the value for the second operation, provide the existing values for the first and the third operations.

Some configuration options cannot be changed using `-modify_current` or `-modify_version`. To change any of the following options, you must create an entirely new SLP version:

<code>-uf</code>	The type of the operation.
<code>-managed</code>	The retention type for the operation: Fixed, Capacity managed, or Expire after copy.
<code>-source</code>	The source of an operation, used primarily in hierarchical SLP configurations.
<code>-dc</code>	The data classification of an existing version.
	The number of operations. You cannot add an operation or remove an operation from the SLP definitions.

See [“Creating a storage lifecycle policy”](#) on page 602.

You cannot instruct an SLP to follow the configuration of a previous version that has been superseded. To revert to the behavior of a previous version, change the definition to match the earlier definition. The change creates a version with the same content as the previous version, but with a new version number.

When changes to storage lifecycle policies become effective

For the changes to become effective for a backlog of jobs, it may be necessary to cancel the applicable jobs.

When the `nbstl` command is used to alter an existing storage lifecycle policy version, those changes may not become effective immediately. The images that are managed by the SLP version that was altered may already belong to a job that is Active or Queued, as seen in the Activity Monitor. Once a job is queued, the characteristics (SLP attributes) are fixed for that job and subsequent changes to the definition have no effect. To make changes effective for a backlog of jobs, cancel the duplication jobs. The storage lifecycle policy manager creates and submits new duplication jobs for those images, using the changes to the configuration.

The following are conditions under which changes to an existing version are not immediately effective:

- Changes to a **Backup** operation have no effect because the backup job is already underway or completed.
- Changes to a **Duplication** operation do not affect the copies that previous duplication jobs created.
- Changes to a **Duplication** operation do not affect the copies that have already been submitted and are currently represented by a duplication job in the Activity Monitor, whether it be Active or Queued. If you want your changes to apply to those active duplication jobs, cancel the applicable duplication jobs. Once the job is canceled, `nbstserv` re-forms and re-submits new duplication jobs for these copies, using the changes to the appropriate version of the SLP.
- Changes to a **Duplication** operation affect the copies that have not yet been created and have not yet been submitted. (That is, they are not yet represented by a duplication job in the Activity Monitor). Your changes become effective for the next duplication session. Whenever `nbstserv` begins a new session, it re-reads the definitions for processing instructions.
- If a duplication job does not complete successfully, unfinished images in the job are submitted as part of a new job. Changes to the version affect the resubmitted job.

Deleting old storage lifecycle policy versions

When a version of a storage lifecycle policy is no longer the active (or most recent) version, the version is subject to deletion. NetBackup automatically deletes the inactive version after all the copies that refer to it have finished processing. When the copies are complete, they are considered SLP-complete.

By default, NetBackup deletes an inactive SLP version after 14 days.

The following parameters in the **SLP settings** host properties apply to version deletion:

- **Cleanup interval** (SLP.CLEANUP_SESSION_INTERVAL)

- **Unused SLP definition version cleanup delay**
(SLP.VERSION_CLEANUP_DELAY)

See [“SLP settings properties”](#) on page 168.

Configuring backups

- [Chapter 20. Creating backup policies](#)
- [Chapter 21. Synthetic backups](#)
- [Chapter 22. Protecting the NetBackup catalog](#)
- [Chapter 23. About the NetBackup database](#)
- [Chapter 24. Managing backup images](#)
- [Chapter 25. Configuring immutability and indelibility of data in NetBackup](#)

Creating backup policies

This chapter includes the following topics:

- [About the Policies utility](#)
- [Planning for policies](#)
- [Creating a backup policy](#)
- [Adding or changing schedules in a policy](#)
- [Changing multiple policies at one time](#)
- [About automanaged policies or storage lifecycle policies](#)
- [Copying or moving policy items to another policy](#)
- [Copying a policy to create a new policy](#)
- [Copying a schedule into the same policy or different policy](#)
- [Deleting schedules, backup selections, or clients from a policy](#)
- [Policy Attributes tab](#)
- [Schedules tab](#)
- [Schedule Attributes tab](#)
- [Start Window tab](#)
- [Excluding days from a schedule](#)
- [Include Dates tab](#)
- [How NetBackup determines which schedule to run next](#)
- [About schedule windows that span midnight](#)

- [How open schedules affect calendar-based and frequency-based schedules](#)
- [About the Clients tab](#)
- [Backup Selections tab](#)
- [Disaster Recovery tab](#)
- [Creating a Vault policy](#)
- [Creating a BigData policy](#)
- [Performing manual backups](#)
- [Active Directory granular backups and recovery](#)

About the Policies utility

Backup policies provide the instructions that NetBackup follows to back up clients. Use the **Policies** utility to provide the following instructions for a backup:

What type of client to back up. See [“Policy Attributes tab”](#) on page 663.

Where to store the backup. See [“Policy Attributes tab”](#) on page 663.

When and how frequently to perform the backup. See [“Schedules tab”](#) on page 730.

Which clients to back up. See [“About the Clients tab”](#) on page 779.

Which client files and directories to back up. See [“Backup Selections tab”](#) on page 782.

Actions in the Policies utility

To perform actions in the Policies utility

- 1 On the left, select **Protection > Policies**.
- 2 Review the following information:

To open a policy: Select the **Policies** tab. Select the link for a policy.

To display specific configuration information: Select the appropriate tab to view the policy configuration areas:

- **Clients** - displays all of the clients that are in the policies.
- **Schedules** - displays all of the schedules that are in the policies.

To display information about all policies on the current primary server: Select the **Policies** tab.

To edit policy attributes for multiple policies:	On the Policies tab: ■ Select the policies that you want to change. ■ Select Edit.
To edit the host properties for a client:	On the Clients tab: ■ Select the client. ■ Select Host properties.
To fetch the resiliency status for a client	On the Clients tab: ■ Select the client. ■ Select Fetch resiliency status.
To perform a manual backup	On the Policies, Clients, or Schedules tab: ■ Select the policies, clients, or schedules for which you want to run a manual backup. ■ Select Manual backup. More details are available in the following topic.

Planning for policies

Policy configuration is flexible enough to meet the various needs of all the clients in a NetBackup environment. To take advantage of this flexibility, take time to plan before starting to configure the policies in the **Policies** utility.

The following table outlines the steps to take to ensure that you get optimal results from your policy configurations.

Table 20-1 Steps for planning policies

Step	Action	Description
Step 1	Gather client information	Gather the following information about each client: ■ The client name. ■ The approximate number of files on each client to be backed up. ■ The typical file size of the files. One client may be a file server that contains a large amount of data while the other clients are workstations. To avoid long backup times, include the file server in one policy and the workstations in another policy. It may be beneficial to create more than one policy for the file server.

Table 20-1 Steps for planning policies (*continued*)

Step	Action	Description
Step 2	Group the clients based on backup requirements	Divide the clients into groups according to the different backup and archive requirements. The groups can be based on the type of work that the clients perform. Clients that are used for similar tasks generally have similar backup requirements. For example, most clients in an engineering department create the same types of files at similar levels of importance. In some instances, create a single policy for each group of clients. In other cases, subdivide the clients and include them in the separate policies that are based on their backup requirements. A backup policy can apply to one or more clients. Every client must be in at least one backup policy so that it can be backed up.
Step 3	Consider the storage requirements	The NetBackup environment may have some special storage requirements that the backup policies must accommodate. The storage unit and volume pool settings apply to all the files that are backed up by a policy. If files have special storage requirements, create separate policies for the files, even if other factors are the same, such as schedules. If it is necessary to keep backups for some files on separate media, create a policy that specifies a unique volume pool for those backups. Then, add the media for that volume pool. See “Windows example of one client in multiple policies” on page 652.
Step 4	Consider the backup schedule	Create additional backup policies if the schedules in one policy do not accommodate all clients and files. Consider the following factors when deciding to create additional policies: ■ Best times for backups to occur. To back up different clients on different schedules may require additional policies with different time schedules. For example, create different policies for night-shift and day-shift clients. ■ How frequently the files change. If some files change more frequently than others, the difference may be enough to warrant creating another policy with a different backup frequency. ■ How long backups need to be retained. Each schedule includes a retention setting that determines how long NetBackup keeps the files that are backed up by the schedule. Because the schedule backs up all the files in the backup selection list, all files should have similar retention requirements. Do not include the files whose full backups must be retained forever, together in a policy where full backups are retained for only four weeks.

Table 20-1 Steps for planning policies (*continued*)

Step	Action	Description
Step 5	Group clients by common attributes	Create separate policies for the clients that require similar policy attribute settings. See “Policy attributes that affect how clients are grouped in policies” on page 653.
Step 6	Maximize multiplexed backups	Create separate policies as necessary to maximize the benefits of multiplexed backups. To maximize drive use, multiplex the slower clients that produce small backups. The higher-performance clients that produce long backups are likely to use drives fully and not benefit from multiplexing. See “Media multiplexing (schedule attribute)” on page 756.
Step 7	Evaluate backup times	Evaluate total backup times for each schedule and further subdivide policies to reduce backup times to an acceptable level. For example, if the backup of <code>/usr</code>, <code>/h001</code>, and <code>/h002/projects</code> on UNIX client1 takes too much time, create a new policy for <code>/h002/projects</code>. In addition to reducing the backup time for each policy, separate policies can reduce the total backup time for the server. NetBackup processes files within a backup selection list in the order they appear in the backup selection list. However, separate policies are processed in parallel if enough drives are available and the Maximum jobs per client host property is set to allow it. See “Global attributes properties” on page 109. The Multiplexing and Allow multiple data streams policy attributes also allow backup policies to be processed in parallel. See “Allow multiple data streams (policy attribute)” on page 697.

See [“About the Policies utility”](#) on page 649.

See [“Policy Attributes tab”](#) on page 663.

Windows example of one client in multiple policies

The following table shows that the files in two different subdirectories on one client can be stored in two different locations.

- Policy1 sends backups of `E:\h002\projects` to DLT storage.
- Policy2 sends backups of `E:\h002\DevExp` and `E:\h002\DesDoc` to DLT storage.

Table 20-2 One client in multiple policies

Policies	Client	Files	Storage
Policy1	client1	C:\ D:\User D:\h001 E:\h002\Projects	DLT
Policy2	client1 client2	E:\h002\DevExp E:\h002\DesDoc	DLT

Policy attributes that affect how clients are grouped in policies

The following table lists the attributes that may determine which clients are grouped in the same policy.

Table 20-3 Policy attributes that affect how clients are grouped in policies

Attribute	Description
Policy Type	Each client must be in a policy of the correct policy type. For example, Windows clients must be in a policy of a MS-Windows policy type. See “Policy type (policy attribute)” on page 664.
Destination	All of the data that the policy generates is sent to the same destination that is indicated in the policy. The data must share the same Data Classification , Policy storage , and Policy volume pool . See “Data classifications (policy attribute)” on page 668. See “Policy storage (policy attribute)” on page 669. See “Policy volume pool (policy attribute)” on page 672.
Job Priority	This attribute determines the priority for the backups of all of the clients in the policy. See “Job priority (policy attribute)” on page 680.
Follow NFS	Select this attribute if a UNIX client has NFS mounted files to be backed up. Consider placing these clients in a separate policy so problems with NFS do not affect the other clients. See “Follow NFS (policy attribute)” on page 682.
Cross mount points	This attribute lets NetBackup cross file system boundaries for all clients in the policy. See “Cross mount points (policy attribute)” on page 685.

Table 20-3 Policy attributes that affect how clients are grouped in policies
(continued)

Attribute	Description
Backup Network Drives	This attribute lets NetBackup back up the files that all clients in the policy store on network drives. (Applies only to the MS-Windows policy type.) See “Backup Network Drives (policy attribute)” on page 683.
Compression	This attribute indicates that all clients in the policy are to compress their backups before they send them to the server. Note that the time to compress can increase backup time and make it unsuitable to use for all clients. Consider creating a different policy for those clients. See “Compression (policy attribute)” on page 689.

About Microsoft DFSR backups and restores

NetBackup protects the databases that are associated with the independent DFSR servers and the DFSR data.

In an environment that has DFSR, two changes occur in NetBackup, as follows:

- To preserve data integrity, the folder or folders that host the Shared Replica DFSR data are excluded automatically by NetBackup from normal file system backups.
- The top-level DFSR shared folders become part of the Shadow Copy Components. Therefore, the data is snapped consistently by Windows Volume Shadow Copy Service (VSS) before each backup.

The VSS writer stops and restarts the DFS Replication service automatically. Schedule the backups to coincide with a period of low activity. (If you stop the replication service manually, Microsoft change journal problems may occur. Specifically, Update Sequence Number (USN) Journal wrap may occur.)

By default all Windows clients are configured for Windows open file backups. The DFSR servers must be configured for this option.

See [“Windows open file backup tab of the Client attributes properties”](#) on page 71.

Refer to [Table 20-4](#) recommendations on how to back up DFSR data.

Table 20-4 Microsoft DFSR backup recommendations

Amount of data	Recommendation
Less than 50 GBs	Configure one policy as follows: ■ Choose the DFSR server host as the client. See “Adding, changing, or deleting clients in a policy” on page 779. ■ Choose ALL_LOCAL_DRIVES as the Directive in the Backup Selections for the policy. The ALL_LOCAL_DRIVES directive includes the Shadow Copy Components:\ automatically. See “Adding backup selections to a policy” on page 783. One policy can back up the data within a reasonable time window.

Table 20-4 Microsoft DFSR backup recommendations (*continued*)

Amount of data	Recommendation
More than 50 GBs	Configure one backup policy for each DFSR server, and in that policy specify only the replication folders. A policy for each host's replication data ensures that the DFSR data is backed up within a reasonable time window. For each DFSR server host, do the following: ■ Create a global exclude list for All Policies and All Schedules. Exclude the following DFSR top-level folder: Shadow Copy Components:\User Data\Distributed File System Replication\DfsrReplicatedFolder The global exclude list ensures that the DFSR components are not backed up accidentally by other backup policies for the client. See “Exclude list properties” on page 94. ■ Create a backup policy for the DFSR data, as follows: ■ For the client, specify the DFSR server host. For the servers that are hosted in a cluster, specify the DFSR cluster name rather than the local host name. See “Adding, changing, or deleting clients in a policy” on page 779. ■ For the Backup Selections for the policy, specify the absolute path to each of the top-level DFSR folders on that host. The following is an example path: Shadow Copy Components:\User Data\Distributed File System Replication\DfsrReplicatedFolders\folder_name <i>Tip:</i> Use the Backup, Archive, and Restore interface to browse the Shadow Copy Components for the DfsrReplicatedFolders folder. The interface shows the path to each DFSR folder that you need to enter as a backup selection. See “Adding backup selections to a policy” on page 783. ■ For the backup policy, create an exception to the exclude list and specify the top-level DFSR directory, as follows: Shadow Copy Components:\User Data\Distributed File System Replication\DfsrReplicatedFolders For the exception Policy, specify the backup policy for the DFSR data. Also specify All Schedules for the Schedule. If DFSR is hosted in a cluster, create the exception for each host in the cluster. The exception ensures that the Shadow Copy Components DFSR paths are included for backup after NetBackup processes the global exclude list. See “Add an exception to the exclude list” on page 96.

During a backup, Windows writes the following event ID messages to the application event log of a DFSR host:

```
Event ID=1102
Severity=Informational
The DFS Replication service has temporarily stopped replication
```

because another application is performing a backup or restore operation. Replication will resume after the backup or restore operation has finished.

Event ID=1104

Severity=Informational

The DFS Replication service successfully restarted replication after a backup or restore operation.

Restores of DFSR

To restore DFSR data, use the NetBackup **Backup, Archive and Restore** client interface to browse the **Shadow Copy Components** for the files or folders to restore, as follows:

Shadow Copy Components:\User Data\Distributed File System
Replication\DfsrReplicatedFolders*folder_name*

When you perform a restore, consider carefully if the restore needs to include the DFSR database, in addition to the DFSR data. A DFSR server maintains a globally unique version number (GVSN) for each DFSR database on each replicated volume. If you restore a DFSR server to an earlier database version, the other servers do not recognize the older version number. Then replication is impossible and stops until the issue is corrected.

More information

https://www.veritas.com/content/support/en_US/article.100038589

Refer to the Microsoft documentation on managing and using DFSR for additional details.

Creating a backup policy

Use the following procedure to create a backup policy.

To create a policy

- 1 In the **NetBackup web UI**, select **Protections > Policies**.
- 2 Click **Add**.
- 3 Enter the policy name.
- 4 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.

Adding or changing schedules in a policy

Change policies only when no backup activity is expected for the affected policies and clients. Make adjustments before backups begin to ensure an orderly transition from one configuration to another.

Changing a policy causes NetBackup to recalculate when the policy is due.

Note: It is not recommended that users modify automanaged policies. If a user begins to modify an automanaged policy, a dialog appears that warns users about the possible consequences.

See [“About automanaged policies or storage lifecycle policies”](#) on page 660.

Use the following procedure to add or change schedules in an existing NetBackup policy.

To add or change schedules in a policy

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Expand the policy name in the middle pane, then select **Schedules**.
- 3 Perform one of the following actions:

Add a schedule

Select **Actions > New > Schedule**.

Copy a schedule and paste it into another policy

- Expand the policy which contains a schedule that you'd like to copy.
- In the right pane, right-click the schedule and select **Copy**.
- Expand the policy where you'd like to paste the schedule.
- In the right pane, right-click anywhere in the schedule area and select **Paste**.

Change an existing schedule Double-click the schedule name.

- 4 Complete the entries in the **Attributes** tab, **Start Window** tab, **Exclude Days** tab, and **Include Dates** tab (when applicable).
 See [“Schedule Attributes tab”](#) on page 731.
 See [“Start Window tab”](#) on page 763.
 See [“Excluding days from a schedule”](#) on page 766.
 See [“Include Dates tab”](#) on page 767.
- 5 Click **OK**.

Changing multiple policies at one time

Use the following procedure to change more than one NetBackup policy at the same time.

Note: It is not recommended that users modify automanaged policies. If a user begins to modify an automanaged policy, a dialog appears that warns users about the possible consequences.

See [“About automanaged policies or storage lifecycle policies”](#) on page 660.

Note: You can change **Attributes**, **Clients**, and **Schedules** of multiple policies at one time. However, you cannot change **Backup Selections** of multiple policies at one time.

To change multiple policies

- 1 Expand **NetBackup Management** and select **Policies**.
- 2 In the middle pane, expand **Summary of All Policies** and select the node that you want to change.
- 3 Select the policies that you want to edit.

Note: You can change different schedules across different policies at one time. If you change multiple schedules of different policy types, the **Attributes** tab is disabled. A tri-state is displayed for an attribute that has different values for multiple policies you selected.

- 4 Select **Edit > Change**.
- 5 Make the desired changes.
See [“Policy Attributes tab”](#) on page 663.
See [“Schedule Attributes tab”](#) on page 731.
See [“About the Clients tab”](#) on page 779.

About automanaged policies or storage lifecycle policies

When a workload administrator subscribes an asset to a protection plan, NetBackup also creates an automanaged policy and automanaged storage lifecycle policy (SLP). In the NetBackup web UI, the name of an automanaged policy or SLP is made up of the protection plan name, then the suffix that contains the GUID. (If you used the custom prefix option when you created the protection plan, then the name of the automanaged policy reflects that prefix instead.) In the NetBackup Administration Console, an automanaged policy or SLP name uses the prefix SLO_ENGINE_MANAGED+.

The NetBackup web UI does not allow users to modify or delete automanaged policies or SLPs. You can only view the details of an automanaged policy or SLP.

Warning: It is not recommended that users modify or delete automanaged policies or storage lifecycle policies using the **NetBackup Administration Console** or the command line. If a user begins to modify or delete an automanaged policy or SLP using the **NetBackup Administration Console**, a dialog appears that warns users about the possible consequences.

- If the user continues to make modifications, they must make sure that the policy or SLP continues to meet the service level objective that is the protection plan defines.
- If the user continues to delete the policy or SLP, they must make sure that the asset is added to another protection plan that meets the service level objective.

Copying or moving policy items to another policy

You can copy or move entire policies, attributes, schedules, clients, and backup selections from one policy to another. The following procedure describes which policy items can be copied or moved.

To copy or move items from one policy to another

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2** In the middle pane, select either the **Attributes**, **Schedules**, **Clients**, or **Backup Selections** for a policy that you want to copy or move.
- 3** In the right pane, based on your selection in the previous step, select the attributes, schedules, clients, or backup selections of a policy that you want to copy or move.
- 4** Do one of the following:

To copy an item	On the Edit menu, click Copy .
To move an item	■ On the Edit menu, click Cut. ■ Click Yes when asked if you want to delete the selected item from the policy.
- 5** In the middle pane, select the policy item to which you want to paste the copied items.
- 6** From the **Edit** menu, click **Paste**.

When you paste items with the same name, NetBackup provides options: To either copy and replace the existing item, or copy but keep the existing item, or to not copy.

Note: If the schedules do not match the policy type, the schedules are not copied. The action is indicated in a dialog box.

When you paste attributes, the existing attributes of the policy are always replaced. Whereas, when you paste backup selections, the backup selection is always copied to the policy, and not replaced.

The copying or moving feature is also applicable to instances and instance groups of Oracle and SQL type of backup policies.

Copying a policy to create a new policy

Use the **Copy to New** option to save time creating policies. This option is especially useful for the policies that contain many of the same policy attributes, schedules, clients, or backup selections.

To copy a policy to create a new one

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2** In the middle pane, select the policy to copy.
- 3** On the **Edit** menu, click **Copy to New**
- 4** In the **Copy a Policy** dialog box, enter the name of the policy that you want to copy. You can indicate a policy other than the one that is selected
- 5** Enter the name for the new policy.
- 6** Click **OK**. The only difference between the new policy and the copied policy is the name.

Copying a schedule into the same policy or different policy

Use the **Copy to New** option to save time creating schedules. Use this option to copy a schedule into the same policy or different policy.

To copy a schedule to create a new one

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2** In the middle pane, expand a policy and select the **Schedules** node that contains the schedule that you want to copy.
- 3** In the right pane, select the schedule that you want to copy.
- 4** On the **Edit** menu, click **Copy to New**
- 5** In the **Copy Schedule** dialog box, enter the name of the new schedule.
- 6** Use the menu to select the name of the policy to which you want to paste the schedule. You can paste the schedule into the same policy or a different policy.
- 7** Click **OK**. The **Change Schedule** dialog box opens for the new schedule.

Deleting schedules, backup selections, or clients from a policy

Use the following procedure to delete schedules, backup selections, or clients from a NetBackup policy.

Note: It is not recommended that users modify or delete automanaged policies. If a user begins to modify or delete an automanaged policy, a dialog appears that warns users about the possible consequences.

See [“About automanaged policies or storage lifecycle policies”](#) on page 660.

To delete a schedule, backup selections, or clients from a policy

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Expand the policy name in the middle pane, and then select **Attributes**, **Schedules**, **Backup Selections**, or **Clients**.
- 3 In the right pane, select the item you want to delete.
- 4 On the **Edit** menu, click **Delete**.
- 5 Click **Yes** when asked if you want to delete the selected item from the policy.


When a client is deleted from the client list, the NetBackup client software is not deleted or uninstalled from the client. Backups for the client can be recovered until the backups expire. Also, when a file is deleted from a backup selection list, the actual file is not deleted from the client.

Policy Attributes tab

Use the policy **Attributes** tab to configure backup settings when you add a new policy or change an existing policy. When you create a policy, you give the policy a name and select a policy type. The policy type you select typically depends on the type of client you want to back up. The number of policy types available varies depending on which NetBackup options are installed. Each policy type has a unique combination of attributes. Not all attributes apply to every policy type. When you select a policy type, the attributes that apply to that policy type are active. The unavailable attributes are grayed out.

[Figure 20-1](#) shows the **Attributes** tab of a NetBackup policy.

Figure 20-1 Policy Attributes tab



The following topics describe the settings on the policy **Attributes** tab.

Policy type (policy attribute)

The **Policy type** attribute determines the purpose of the policy. Select a policy type from the list. The policy type you select typically depends on the type of client to be backed up. Some policy types are not used for client backups. For example, **NBU-Catalog** is not used for client backups.

The list of policy types changes depending on the NetBackup options that have been installed. Each policy type offers a unique combination of attributes. When you select a policy type, only the attributes that apply to that policy type are active.

You can change the policy type of an existing policy. However, the schedules for the policy may become invalid. If the schedules become invalid, NetBackup displays

a warning message and then deletes the invalid schedules or changes the schedules to an equivalent type.

When you change the policy type of an existing policy, other selections or features of the policy may become invalid as well.

[Table 20-5](#) describes the types of NetBackup policies that are available in the NetBackup web UI.

Table 20-5 NetBackup policy types

Policy type	Description
BigData	Use as a policy type to schedule and run a job for big data applications. For example, Hadoop Distributed File System (HDFS) or Nutanix Acropolis Hypervisor. This policy type requires the Enterprise Client license. See “Creating a BigData policy” on page 832. For information about the plug-ins that NetBackup supports, and which are available to download and install, see: https://www.veritas.com/content/support/en_US/article.100040155
Cloud	This policy type is only available in the NetBackup web UI. Use for policies that back up in-cloud workloads. For example, assets in Amazon Web Services (AWS), Azure, Microsoft Azure and Azure Stack Hub, Google Cloud Platform (GCP), and Oracle Cloud Infrastructure (OCI). For a complete list of compatible in-cloud workloads, refer to the NetBackup Hardware Compatibility List (HCL).
Cloud-Object-Store	This policy type is only available in the NetBackup web UI. Use for policies that back up Cloud object stores. For example, Azure Blob Storage, AWS S3, Google Cloud Storage (GCS), and Hitachi Cloud Platform. For a complete list of compatible object stores, refer to the NetBackup Hardware Compatibility List (HCL).
DataTools-SQL-BackTrack (UNIX only)	This policy type is only available in the NetBackup Administration Console. Use for the policies that contain only clients with the NetBackup for SQL-BackTrack agent. For information on setting up this policy type, see the guide for this option.
DataStore	This policy type is used for database applications that NetBackup uses the Open Backup Services (XBSA) for backup or archive purposes.
DB2	Use for the policies that contain only clients with the NetBackup for DB2 agent. For information on setting up this policy type, see the guide for this option.
Enterprise-Vault	Use as a policy type to schedule and run an Enterprise Vault job.

Table 20-5 NetBackup policy types (*continued*)

Policy type	Description
FlashBackup (UNIX only)	Combines the speed of raw-partition backups with the ability to restore individual files. This policy type applies to UNIX clients only. Requires the Enterprise Client license. For information on setting up this type of policy, see the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide.
FlashBackup-Windows (Windows only)	Combines the speed of raw-partition backups with the ability to restore individual files. This policy type applies to Windows clients only. Requires the Enterprise Client license. For information on setting up this type of policy, see the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide.
Hyper-V	To back up the virtual machines that reside on Windows Hyper-V servers, by means of on-host or off-host backups. This policy type requires the Enterprise Client license. Users can upgrade pertinent policies to the Hyper-V policy type using one of the following methods: - Change the policy type in the NetBackup web UI or the NetBackup Administration Console. In the NetBackup Administration Console, change the type for only one policy at a time. - Change the policy type for multiple policies at one time using the <code>bpplconvert</code> script that is located in the following location: On Windows: <code>install_path\NetBackup\bin\goodies</code> On UNIX: <code>usr/openv/netbackup/bin/goodies</code> For information on setting up this type of policy, see the NetBackup for Hyper-V Guide.
Kubernetes	Use to protect Kubernetes clusters.
Informix-On-BAR (UNIX only)	Use for the policies that contain only clients with the NetBackup for Informix agent. For information on setting up this policy type, see the guide for this option.
Lotus-Notes	Use for the policies that contain only clients with the NetBackup for Lotus Notes agent. For information on setting up this policy type, see the guide for this option.
MongoDB Ops Manager	Use for the policies that contain only clients with the NetBackup for MongoDB Ops Manager agent. For information on setting up this policy type, see the guide for this option.
MS-Exchange-Server (Windows only)	Use for the policies that contain only clients with the NetBackup for Exchange Server agent. For information on setting up this policy type, see the guide for this option. This policy type applies to Windows clients only.

Table 20-5 NetBackup policy types (*continued*)

Policy type	Description
MS-SharePoint (Windows only)	Use to configure a policy NetBackup for SharePoint Server. This policy type applies to Windows clients only.
MS-SQL-Server	Use for the policies that contain only clients with the NetBackup for SQL Server agent. For information on setting up this policy type, see the guide for this option.
MS-Windows	Use for the policies that contain only Windows clients of supported Windows operating system levels. Standard and MS-Windows policy types are the only policy types that support the following options: ■ Checkpoint restart for backups or restores See “Take checkpoints every __ minutes (policy attribute)” on page 674. ■ Synthetic backups See “Policy type (policy attribute)” on page 664. ■ Collect disaster recovery information for Bare Metal Restore See “Collect disaster recovery information for Bare Metal Restore (policy attribute)” on page 693.
NAS-Data-Protection	Use the policy to configure dynamic data streaming for NAS workloads. For information on setting up this type of policy, see the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide .
NBU-Catalog	Use for catalog backup jobs.
NDMP	Use for the policies that contain only clients with the NetBackup NDMP option. This policy type is available only when the NetBackup NDMP is installed and licensed. For information on setting up this policy type, see the guide for this option.
Nutanix-AHV	NetBackup web UI supports policy to backup Nutanix-AHV assets. On the NetBackup web UI, click Protection > Policies . Click Add and select Nutanix-AHV as the Policy type . Note: This policy type requires the Enterprise Client license and only available in web UI. For more details, see <i>NetBackup for Nutanix-AHV Administrator Guide > Protecting AHV virtual machines > Backup AHV VMs or intelligent groups using Policy</i> section.
Oracle	Use for the policies that contain only clients with the NetBackup for Oracle agent. For information on setting up this policy type, see the guide for this option.
SAP	Use for the policies that contain only clients with the NetBackup SAP Agent. For information on setting up this policy type, see the guide for this option.

Table 20-5 NetBackup policy types (*continued*)

Policy type	Description
Standard	Use for UNIX clients (including supported Mac clients), except for those clients that are covered by specific products, such as Oracle. Standard and MS-Windows policy types are the only policy types that support the following options: ■ Checkpoint restart for backups or restores ■ Synthetic backups ■ Collect disaster recovery information for Bare Metal Restore
Sybase	Use for the policies that contain only clients with the NetBackup for Sybase agent. For information on setting up this policy type, see the guide for this option.
Universal-Share	Use as a policy type to schedule and run a job that protects data in a universal share. The universal share must be created on a specified MSDP storage server using NetBackup storage APIs or the NetBackup web UI .
Vault	This policy type is only available in the NetBackup Administration Console. Use as a policy type to schedule and run a Vault job. This policy type is available only when Vault is licensed.
VMware	For backup of any virtual machines that reside on VMware vSphere servers, by means of off-host backups. This policy type requires the Enterprise Client license. Users can upgrade pertinent policies to the VMware policy type using one of the following methods: ■ Change the policy type in the NetBackup web UI or the NetBackup Administration Console. In the NetBackup Administration Console, change the type for only one policy at a time. ■ Change the policy type for multiple policies at one time using the <code>bpplconvert</code> script that is located in the following location: On Windows: <code>install_path\NetBackup\bin\goodies</code> On UNIX: <code>usr/openv/netbackup/bin/goodies</code> For information on setting up this type of policy, see the NetBackup for VMware Guide.

For information about off-host backups, see the [NetBackup NAS Administrator's Guide](#) and the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).

Data classifications (policy attribute)

The **Data Classification** attribute specifies the classification of the storage lifecycle policy that stores the backup. For example, a backup with a gold classification must

go to a storage unit with a gold data classification. By default, NetBackup provides four data classifications: platinum, gold, silver, and bronze.

This attribute is optional and applies only when the backup is to be written to a storage lifecycle policy. If the list displays **No data classification**, the policy uses the storage selection that is displayed in the **Policy storage** list. If a data classification is selected, all the images that the policy creates are tagged with the classification ID.

See [“Storage lifecycle policy settings”](#) on page 629.

See [“Data Classification properties”](#) on page 83.

See [“Adding a Data Classification”](#) on page 84.

See [“About storage lifecycle policies”](#) on page 601.

Policy storage (policy attribute)

The **Policy storage** attribute specifies the storage destination for the policy’s data. Select a specific storage unit, storage lifecycle policy, or storage unit group from the list.

When NetBackup looks for an available storage unit, it selects the first storage unit that meets the following requirements:

- The storage unit must not be designated as **On demand only**.
- The storage unit must have available drives.
- The storage unit must have media available in the required volume pool.

However, NetBackup makes an exception when a client is also a media server with locally-attached storage units. In that case, NetBackup selects the locally-attached storage units first.

See [“About staging backups”](#) on page 576.

Storage unit

Select the name of the storage unit that is to be the storage destination for the policy data. It can be disk or tape-based.

If it is configured to do so, the storage unit determines which type of disk staging is used for the policy.

See [“Creating a storage unit”](#) on page 548.

Storage lifecycle policy	Select the name of the storage lifecycle policy that is to be the storage destination for the policy data. The drop-down list includes only those lifecycles that have the same data classification as the policy. For example, gold backup images cannot be sent to a silver storage lifecycle. Images that belong to a specific data classification cannot be sent to a storage lifecycle that lacks a classification. Data classification is optional. See “Global attributes properties” on page 109. If it is configured to do so, the storage lifecycle policy determines which type of disk staging is used for the policy. If setting up snapshot replication with Replication Director, select a storage lifecycle policy that contains a snapshot-capable storage unit. See “About storage lifecycle policies” on page 601.
Storage unit group	Select the name of the storage unit group that is to be the storage destination for the policy data. See “About storage unit groups” on page 590.
Any Available	If Any Available is selected, NetBackup tries to store data on locally-attached storage units first. To force NetBackup to use only a locally-attached drive, select Must use local drive in the General Server properties. If a local device is not found or Must use local drive is not selected, NetBackup tries to find an available storage unit alphabetically. NetBackup does not select a null_stu storage unit if Any Available is selected. A null_stu storage unit is created only when Cohesity Technical Support uses the NullOST plug-in to identify and isolate data transfer bottlenecks.

Figure 20-2 Icons indicate type of storage

-  Storage unit (tape device)
-  Storage unit (disk)
-  Storage unit group
-  Storage lifecycle policy

Note: If different storage is selected for the **Override policy storage** option on the **Schedule Attributes** tab, that selection overrides the **Policy storage** attribute.

See [“Override policy storage \(schedule attribute\)”](#) on page 751.

See [“Considerations for selecting a destination for Policy storage”](#) on page 671.

Considerations for selecting a destination for Policy storage

Consider the following scenarios before selecting a destination from the **Policy storage** list on the policy **Attributes** tab.

Table 20-6

Scenario	Action
The site contains one storage unit, or there is no storage unit preference.	Do one of the following: ■ Specify Any Available for the Policy storage attribute. ■ Do not specify a storage unit at the schedule level. See “Override policy storage (schedule attribute)” on page 751. ■ Do not set all storage units to On demand only. NetBackup may not find an available storage unit for the backups. See “Editing storage unit settings ” on page 551. See “On demand only storage unit setting” on page 568.
A specific storage unit is designated but the unit is unavailable.	Consider changing the destination to Any Available since backups cannot run for those policies and the schedules that require the unit.
Any Available is selected.	Be aware that any basic disk storage unit that is not assigned to a storage group is considered available for disk spanning. See “Media properties” on page 119.

Table 20-6 (continued)

Scenario	Action
You want to limit the storage units available to a policy.	Do one of the following: ■ Select a storage unit group that contains only the units you want the policy to use. ■ Limit the storage units by doing the following: ■ Create a volume pool that contains the volumes that are available only to the specific storage units. Disable Scratch pool for the volume pool. If Scratch pool is enabled, any storage unit has access to the volumes in the volume pool. See “Add a volume pool” on page 521. See “Edit or delete a volume pool” on page 521. See “About scratch volume pools” on page 465. ■ In the policy, set Policy volume pool to the volume pool that is defined in the previous step. ■ For all policies, set Policy storage attribute to Any Available. ■ If the policy specifies a storage unit group, set the storage units within the group to On demand only to satisfy the policy requirement. See “Editing storage unit settings ” on page 551. See “On demand only storage unit setting” on page 568.

Policy volume pool (policy attribute)

The **Policy volume pool** attribute specifies the default volume pool where the backups for the policy are stored. A volume pool is a set of media that is grouped for use by a single application. The volume pool is protected from access by other applications and users.

The available volume pools appear on the list. Whenever a new volume is required, it is allocated from the volume pool indicated.

If you select a volume pool on the **Schedule** tab, that selection overrides the **Policy volume pool** selection on the **Attributes** tab.

See [“Override policy storage \(schedule attribute\)”](#) on page 751.

See [“Example of overriding the policy volume pool”](#) on page 673.

The following table describes the default volume pools that NetBackup defines.

Table 20-7 Default volume pools defined by NetBackup

Volume pool	Description
None	The default pool for applications, other than NetBackup.

Table 20-7 Default volume pools defined by NetBackup (*continued*)

Volume pool	Description
DataStore	The default pool for DataStore.
NetBackup	Unless otherwise specified in the policy, all backups use media from the NetBackup pool. One exception is the NBU-Catalog policy type.
CatalogBackup	This pool is selected by default for the NBU-Catalog policy type. It is used exclusively for catalog backups. Catalogs are directed to a single, dedicated pool to facilitate faster catalog restores.

The following table describes the additional volume pools that are useful to create.

Table 20-8 Additional volume pools

Volume pool	Description
Scratch volume pool	Allows NetBackup to automatically transfer volumes when another volume pool does not have media available.
Auto volume pool	Used by automatic backups.
User volume pool	Used by user backups.

Media is assigned to the volume pools for Media Manager storage devices. Disk-type storage devices are not allocated to a volume pool.

See [“About NetBackup volume pools”](#) on page 463.

See [“Add a volume pool”](#) on page 521.

See [“Edit or delete a volume pool”](#) on page 521.

See [“About scratch volume pools”](#) on page 465.

Example of overriding the policy volume pool

The following example shows how to override the policy volume pool from the policy **Schedule** tab. In this example, we change a policy named *Backup-Archive*. Until now, all schedules in the policy have used the *Backups* volume pool. Change the policy so that the user-archive schedule uses the *Archive* pool instead.

To override the Policy volume pool attribute

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**
- 2** In the left pane, select the *Backup-Archive* policy and on the **Edit** menu, click **Change**.
- 3** In the policy **Attributes** tab, on the **Policy volume pool** list, select *Backups*.
- 4** Click the **Schedules** tab.
- 5** Select the schedules that use the *Backups* volume pool, and click **Properties**.
- 6** Make sure that **Override policy volume pool** is unchecked, and click **OK** to save the change in the schedule.
- 7** Select the user-archive schedule that you want assigned to the *Archive* volume pool, and click **Properties**.
- 8** Check **Override policy volume pool**.
- 9** Underneath the check box, select *Archive* from the list.
- 10** Click **OK** to save the change in the schedule.
- 11** Click **OK** to save the change in the policy.

Take checkpoints every __ minutes (policy attribute)

By taking checkpoints during a backup, you can save time if the backup fails. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the beginning of the last checkpoint. This action is often quicker rather than restarting the entire job.

The checkpoint frequency indicates how often NetBackup takes a checkpoint during a backup. The default is 15 minutes. The administrator determines checkpoint frequency on a policy-by-policy basis. When you select the checkpoint frequency, balance the loss of performance due to frequent checkpoints with the possible time lost when failed backups restart. If the frequency of checkpoints affects performance, increase the time between checkpoints.

Checkpoints are saved at file boundaries and point to the next file in the list. Checkpoint restart is only available for the **MS-Windows**, **NAS-Data-Protection**, or **Standard** policy types. Check **Take checkpoints every __ minutes** to enable a checkpoint restart. When the box is checked, NetBackup takes checkpoints during a backup job at the frequency you specify. If the box is not checked, no checkpoints are taken and a failed backup restarts from the beginning of the job. Checkpoint restart can also be used for restore jobs.

See [“Checkpoint restart for restore jobs”](#) on page 677.

The **Global Attributes** property, **Schedule backup attempts**, indicates the number of times that NetBackup tries to restart a failed backup.

See [“Global attributes properties”](#) on page 109.

Note: Checkpoints are saved at file boundaries and point to the next file in the list to be backed up. Checkpoints cannot occur in the middle of a file. After the file is backed up, the checkpoint is saved.

Note: Checkpoints are not taken for a user-archive backup. If a user-archive backup resumes, it restarts from the beginning.

In the following situations, NetBackup starts a new job instead of resuming an incomplete job:

- If a new job is due to run. Or for calendar-based scheduling, another run day has arrived.
- If the time since the last incomplete backup was longer than the shortest frequency in any schedule for the policy.
- If the time indicated by the Clean-up property, **Move backup job from incomplete state to done state**, has passed.

The following table describes the level of support for various policy attributes, storage, and clients for checkpoint restart. For an agent or option not listed, refer to the manual for that agent or option.

Table 20-9 Support for checkpoint restart

Item	Description
Basic disk staging	Checkpoint restart is supported for Stage I. Checkpoint restart is not supported for Stage II. See “About basic disk staging” on page 577. See “About staging backups” on page 576.

Table 20-9 Support for checkpoint restart (*continued*)

Item	Description
MS-Windows (policy type)	The following pertain to Windows clients: ■ Checkpoint restart is not supported for the backup selections that have a UNC path. ■ No checkpoints are taken during a system state backup. ■ No checkpoints are taken during a Windows disk image (raw) backup. ■ No checkpoints are taken for the remainder of the backup after NetBackup encounters Single-instance Store (SIS). When an incremental backup resumes and completes successfully, the archive bits are cleared for the files that were backed up after the job resumes. However, the archive bits are not cleared for the files that were backed up before the resume. Since the archive bits remain, the files that were backed up before the backup resumes are backed up again during the next incremental backup.
Multiple copies (schedule attribute)	Checkpoint restart is supported for the policies that are configured to create multiple backup copies. See “Multiple copies (schedule attribute)” on page 747. The last failed copy that contains a checkpoint can be resumed. Another copy must be configured to allow other copies to continue the job if the copy fails and subsequent checkpoints occur.
Snapshot Client (policy attribute)	Checkpoint restart is supported for use with local or alternate client backups. However, the following policy attributes are not supported: ■ Block-level incremental backups ■ Instant Recovery backup See “Snapshot Client and Replication Director (policy attributes)” on page 728.
Standard (policy type)	Checkpoint restart is supported for UNIX clients.
Synthetic backups (schedule attribute)	Checkpoint restart is not supported. See “Synthetic backup (schedule attribute)” on page 741.

Table 20-9 Support for checkpoint restart (*continued*)

Item	Description
NAS-Data-Protection (policy type)	The checkpoint restart behavior is applicable at a volume level. You can suspend or resume backup jobs for a volume. The checkpoint interval that is configured for the policy is applicable when a child backup job starts to backup the volume content. 1 Manually suspend and resume jobs: If you suspend a backup from snapshot job (parent job or child job), the entire job hierarchy for that volume goes into a suspended state. When a backup from snapshot job is resumed (parent job or child job), the entire job hierarchy for that volume goes into an active state. 2 Incomplete state: If a child backup from a snapshot job fails due to any reason, the entire job hierarchy for that volume goes into an incomplete state. You can resume incomplete backup jobs after rectifying the failure condition. When you resume any backup from a snapshot job (parent job or child job), the entire job hierarchy for that volume goes into an active state. For more information about the NAS-Data-Protection policy, see the NetBackup NAS Administrator's Guide.

Checkpoint restart for restore jobs

Checkpoint restart for restore jobs saves time by letting NetBackup resume a failed restore job. The job resumes automatically from the start of the file that was last checkpointed rather than starting from the beginning of the entire restore job. NetBackup automatically takes checkpoints once every minute during a restore job.

The following host properties affect checkpoint restart for restore jobs.

Move restore job from incomplete state to done state	This Clean-up host property indicates the number of days that a failed restore job can remain in an Incomplete state. See “Clean up properties” on page 61.
Restore retries	This Universal Setting host property specifies the number of attempts that a client has to restore after a failure. See “Universal settings properties” on page 178.

Checkpoint restart for restore jobs has the following limitations:

- The restore restarts at the beginning of the last checkpointed file, not within the file.
- Only the backups that are created using **MS-Windows** or **Standard** policy types are supported.
- Third Party Copy and the Media Server Copy images that use **Standard** policy types are supported. However, they cannot be suspended or resumed if the backup image has changed blocks.

A NetBackup administrator can choose to suspend a checkpointed restore job and resume the job at a later time. For example, while an administrator runs a restore job for several hours, the administrator receives a request for a second restore. The request is of a higher priority and requires the resources in use by the first job. The administrator can suspend the first job, start the second restore job and let it complete. The administrator can then resume the first job from the Activity Monitor and let the job complete.

Consider a situation in which a checkpointed restore that has no end date is suspended and then resumed. If a new backup occurs before the resume is initiated, the files from the new backup are included in the restore. For example, a user request the restore of a directory. The restore begins, but is suspended. The request is resumed the next day after another backup of the directory is performed. The files that are restored are from the latest backup.


See [“Take checkpoints every __ minutes \(policy attribute\)”](#) on page 674.

Limit jobs per policy (policy attribute)

The **Limit jobs per policy** attribute limits the number of jobs that NetBackup performs concurrently when the policy is run. By default, the box is unchecked, and NetBackup performs an unlimited number of backup jobs concurrently. Other resource settings can limit the number of jobs.

A configuration can contain enough devices so that the number of concurrent backups affects performance. To specify a lower limit, check the box and specify a value from 1 to 999.

Figure 20-3 Limit jobs per policy attribute



Client A and Client B backups can occur concurrently and back up to different devices

This attribute operates differently for VMware policies, depending on how the policy selects virtual machines.

For more information, see the [NetBackup for VMware Administrator's Guide](#).

[Table 20-10](#) describes the factors that affect the number of concurrent backup jobs that NetBackup can perform.

Table 20-10 Factors affecting the number of concurrent backup jobs

Item	Description
Jobs from different policies	The limit does not apply to concurrent jobs if the jobs are from different policies. For example, if three policies limit concurrent jobs to two, NetBackup can start two jobs from each policy. A total of six policies can be running at one time in this situation.
Multiplexing	If multiplexing is used, set the limit high enough to support the specified level of multiplexing. Lower values can limit multiplexing within a policy if jobs from different schedules exist within the policy. For example, the limit is set to two and an incremental backup schedule is due to run for four clients. Only two clients are backed up at one time, regardless of the multiplexing settings.
Network load	The available bandwidth of the network determines how many backups can occur concurrently. If you encounter loading problems, consider multiple networks for backups. Or, configure the backup policy to use the Compression attribute. See "Compression (policy attribute)" on page 689. When the client that is backed up is also a server, it is a special case. In this instance, the network load is not a factor because the network is not used. However, the load on the client and server is still a factor.

Table 20-10 Factors affecting the number of concurrent backup jobs
(continued)

Item	Description
Number of storage devices available and multiplexing limits	To process more than one backup job at a time, the configuration must include one of the following: ■ Multiple storage units. ■ A storage unit with enough drives to perform more than one backup at a time. ■ Storage units that are configured to multiplex. With removable media devices such as tape drives, the number of concurrent jobs depends on the total number of drives in the storage units. With disk storage, the storage device is defined as a file path and the available disk space determines how many paths are possible.
Parent job and children jobs	Parent jobs do not count toward the limit. Only the children jobs count toward the limit. The following are examples of the jobs that produce a parent job and children jobs: ■ Multistreamed jobs ■ Catalog backups ■ Snapshot Client snapshots ■ Bare Metal Restore jobs See “About the Jobs tab” on page 1009. This attribute operates differently for VMware policies, depending on how the policy selects virtual machines. For more information, see the NetBackup for VMware Administrator's Guide.
Server speed	Too many concurrent backups interfere with the performance of the server. The best number depends on the hardware, operating system, and applications that are running.

Job priority (policy attribute)

The **Job priority** attribute specifies the priority that a policy has as it competes with other policies for resources. Enter a value from 0 to 99999. The higher the number, the greater the priority of the job. NetBackup assigns the first available resource to the policy with the highest priority.

In the **Default Job Priorities** host properties, you can set a job priority default for a job type.

See [“Default job priorities properties”](#) on page 85.

Media Owner (policy attribute)

The **Media Owner** attribute specifies which media server or server group should own the media that backup images for this policy are written to.

This attribute is active under the following conditions:

- A Media Manager storage unit is used.
- The **Policy storage** attribute is set to **Any Available**

You can specify the following for the **Media Owner**:

Any (default)	Allows NetBackup to select the media owner. NetBackup selects a media server or a server group (if one is configured).
None	Specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	Allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured appear in the list.

See [“About media sharing”](#) on page 525.

See [“Add a server group”](#) on page 374.

Go into effect at (policy attribute)

The **Go into effect at** attribute specifies when the policy can begin to schedule backups. For example, if today is Monday and you enter Wednesday at 12:00 A.M., the policy does not run until that time or later. Use this attribute to configure a series of policies in advance of when the policies need to become active.

To activate the policy, check **Go into effect at**. The policy must be active for NetBackup to use the policy.

To deactivate a policy, uncheck the box. Inactive policies appear are unavailable in the **NetBackup Administration Console**. Inactive policies remain on the list of policies in the left pane of the **NetBackup Administration Console**. To resume backups, recheck the box. Make sure that the date and time are set to the time that you want to resume backups.

If the schedule is to be used for a catalog archive, the policy must not be active. Clear the check box to deactivate the policy.

See [“Creating a catalog archiving policy”](#) on page 884.

Follow NFS (policy attribute)

The **Follow NFS** (Network File System) attribute specifies whether NetBackup is to back up or archive any NFS-mounted files. These files are named in the backup selection list or by the user, in the case of a user backup or archive. Uncheck the box to prevent the backup or archive of NFS-mounted files.

Note: This attribute applies only to UNIX clients in certain policy types. NetBackup allows it to be selected in those instances only.

This attribute eliminates the need to locate and log on to the systems where the files reside. If the files are mounted on the NetBackup client, you can back up, archive, and restore them by working from the NetBackup client. You must have the necessary permissions on the NFS mount. Use this capability to back up the systems that the NetBackup client software does not support.

Generally, do not back up NetBackup clients over NFS. Back up and archive files on the NFS server where the files physically reside. NFS backups have lower performance and sometimes encounter problems. If **Follow NFS** is selected, you may want to use the policy only for the files and clients that are backed up or archived over NFS.

Note: If **Follow NFS** is not selected, the backup process reads the client's mount table and evaluates each item in the table. NetBackup resolves any links to the true path. NetBackup must resolve the links so it can accurately avoid backing up any files that reside on NFS-mounted file systems.

If NetBackup cannot access a Network File System when it evaluates the mount table, it assumes that the file system is unavailable. (The default time to access the file system is five seconds.) To change the default, change the UNIX primary server host property, `NFS_ACCESS_TIMEOUT` in the `usr/opensv/netbackup/bp.conf` file.

Note: NetBackup specifically excludes mapped directories even if **Follow NFS** and **Cross mount points** are enabled. To back up mapped directories, include the directories in the file list.

Consider the following before enabling this attribute:

Table 20-11 Issues that affect Follow NFS

Item	Description
Cross mount points (policy attribute)	The behavior of Follow NFS can vary depending on how it is used in combination with Cross mount points. See “Examples of using Cross mount points and Follow NFS in combination” on page 687. See “Cross mount points (policy attribute)” on page 685.
Raw partitions	This attribute has no effect on raw partitions. The Network File Systems that are mounted in a raw partition are not backed up. Nor can you back up raw partitions from other computers that use NFS mounts to access the raw partitions. The devices are not accessible on other computers through NFS. Note: NetBackup does not support raw partition backups on unformatted partitions.
Automounted directories	This attribute causes files in automounted file systems to be backed up. Automounted directories can be excluded to allow the backup of other NFS mounts. To do so, add an entry for the automounter’s mount directory to the exclude list on the client.

Backup Network Drives (policy attribute)

The **Backup Network Drives** attribute is for use on single user systems, Win95, Win98, and ME. These operating systems are not supported with this version of NetBackup. For a computer that is not a NetBackup client, the preferred method for backing up data is to use UNC paths. UNC paths are more precise and indicate exactly what should be backed up.

When you use **Backup Network Drives** or UNC paths, the network drives must be available to the service account that the NetBackup Client service logs into at startup. By default, the startup account is set to System. You must change this account on each Windows client that is backed up that contains data that is shared from another computer.

This attribute must be enabled for the policies that back up to CD ROM drives. For scheduled backups, the file list must indicate at least the first level of folders to be backed up. For example, `D:\Folder1` instead of only `D:\`

Note: Mapped drive letters cannot be backed up. Drive letters do not appear in the **Backup, Archive, and Restore** console when backups are browsed.

Example of using UNC paths to back up a shared folder

The following example gives the steps for backing up a shared folder using a UNC path. The procedure backs up the folder *TestData* on *win_PC* through *win_client*. Consult the following descriptions before you review the example.

<i>primary1</i>	NetBackup primary server
<i>win_client</i>	Windows NetBackup client
<i>win_PC</i>	Windows computer (not necessarily a NetBackup client)
<i>TestData</i>	A shared folder on <i>win_PC</i>

Table 20-12 Using UNC paths to back up a shared folder on *win_PC*

Step	Action	Description
Step 1	Create a policy	On <i>primary1</i> create a policy for <i>win_client</i> .
Step 2	Add the folder name to the policy	Add <code>\\win_PC\TestData</code> to the file list of the policy. This step is not necessary if the policy is only used for user-directed backups.
Step 3	Configure the NetBackup Client Service	Perform the following actions: - On <i>win_client</i>, change the NetBackup Client Service to Start Up or Log On with the same account as the user that performs the backup. This user account must have read permissions for the share that is to be backed up. The account must have write permission to perform restores. - Stop and start the NetBackup Client Service so the new account takes effect. See “Configuring the NetBackup Client Service” on page 1052.
Step 4	Perform a backup	Backups run as scheduled or when a manual backup is performed. See “Performing manual backups” on page 834.

Example of using Backup Network Drives (policy attribute) to back up a shared folder

The following example gives the steps for backing up a shared folder using the **Backup Network Drives** policy attribute. The procedure backs up the folder *share* on *win_PC* through *win_client*. Consult the following descriptions before you review the example.

<i>primary1</i>	NetBackup primary server
<i>win_client</i>	Windows NetBackup client

win_PC Windows computer (not necessarily a NetBackup client)

share A shared folder on *win_PC*

Table 20-13 Using Backup Network Drives to back up a shared folder on *win_PC*

Step	Action	Description
Step 1	Create a policy	On <i>primary1</i> create a policy for <i>win_client</i> , and check Backup network drives in the policy attributes tab.
Step 2	Configure the NetBackup Client Service	Perform the following actions: ■ On <i>win_client</i>, change the NetBackup Client Service to Start Up or Log On with the same account as the user that performs the backup. This user account must have read permissions for the share that is to be backed up. The account must have write permission to perform restores. ■ Stop and start the NetBackup Client Service so the new account takes effect. See “Configuring the NetBackup Client Service” on page 1052.
Step 3	Create a batch file	Create a batch file <code>bpstart_notify.bat</code> that does the following: ■ Maps a drive on <i>win_client</i> to <code>\\win_PC\share</code>. ■ Includes the following command (where X: is the mapped drive letter): <pre>net use X: \\win_PC\share</pre>
Step 4	Perform a backup	Backups run as scheduled or when a manual backup is performed. See “Performing manual backups” on page 834.

Cross mount points (policy attribute)

The **Cross mount points** attribute controls whether NetBackup crosses file system boundaries to back up or archive all files and directories in the selected path. For example, if root (/) is specified as the file path on a UNIX system, NetBackup backs up root (/) and all files and directories under root in the tree.

When this attribute is disabled, only the files that are in the same file system as the selected file path are backed up. By disabling, you also prohibit NetBackup from crossing mount points to back up root (/) without backing up all the file systems that are mounted on root. (For example, `/usr` and `/home`.)

In some cases, consider creating separate policies for the backups that cross mount points and those that do not. For example, in one policy, disable **Cross mount**

points and include `root (/)` in the backup selection list. As a result, only the root file system is backed up, and not the file systems that are mounted on it. In another policy, enable **Cross mount points** and include `root (/)` in the backup selection list. As a result, all the data on the client is backed up.

Note: NetBackup specifically excludes mapped directories even if **Follow NFS** and **Cross mount points** are enabled. To back up mapped directories, include the directories in the file list.

The following table lists items to consider when you use this policy attribute.

Table 20-14 Considerations for Cross mount points (policy attribute)

Item	Description
Follow NFS (policy attribute)	The behavior of Cross mount points can vary depending on how it is used in combination with Follow NFS. See “Examples of using Cross mount points and Follow NFS in combination” on page 687. See “Follow NFS (policy attribute)” on page 682.
Backup selection entries	The following backup selection entries behave in the following manner for Windows and UNIX computers when the Cross mount points attribute is used: ■ <code>/</code> Valid for UNIX clients. For Windows clients, the forward slash expands to <code>ALL_LOCAL_DRIVES</code>. ■ <code>:\</code> Valid for Windows clients. For UNIX clients, this entry creates a status 69 (Invalid filelist specification). ■ <code>*:\</code> Valid for Windows clients. For UNIX clients, this entry creates a status 69 (Invalid filelist specification).
UNIX raw partitions	This attribute has no effect on UNIX raw partitions. If a raw partition is the root partition and contains mount points for other file systems, the other file systems are not backed up when Cross mount points is enabled.

Table 20-14 Considerations for Cross mount points (policy attribute)
(continued)

Item	Description
ALL_LOCAL_DRIVES directive	Do not use Cross mount points in policies on UNIX computers where you use the ALL_LOCAL_DRIVES directive in the backup selection list. Enabling Cross mount points can cause multiple backups of mounted volumes. If you require the backup to traverse file system boundaries, do not use the ALL_LOCAL_DRIVES backup selection directive on UNIX clients. Instead, specify a forward slash (/) within the policy backup selection list and ensure that Cross mount points is selected in the policy Attributes.
Mount points to disk storage	Do not cross mount points to back up a media server that uses mount points to any disk storage that contains backup images. If the policy crosses mount points, the NetBackup backup images that reside on that disk storage are backed up. The NetBackup disk storage unit type uses mount points for disk storage.

Examples of using Cross mount points and Follow NFS in combination

By using **Cross mount points** and **Follow NFS** in combination, you can get a variety of results. [Table 20-15](#) summarizes the possible results.

Table 20-15 Results of using Cross mount point and Follow NFS in combination

Cross mount points	Follow NFS	Result
Disabled	Disabled	No crossing of mount points (default).
Disabled	Enabled	Back up NFS files if the file path is (or is part of) an NFS mount.
Enabled	Disabled	Cross local mount points but not NFS mounts.
Enabled	Enabled	Follow the specified path across mount points to back up files and directories (including NFS), regardless of the file system where they reside.

Note: NetBackup specifically excludes mapped directories even if **Follow NFS** and **Cross mount points** are enabled. To back up mapped directories, include the directories in the file list.

Example 1 and Example 2 assume that the client disks are partitioned as shown in [Figure 20-4](#).

Figure 20-4 Example configuration of client disks




Table 20-16 Legend

Disks	Description
d1	Contains /(root), /usr, and /home in separate partitions.
d2	Contains a file system named /home/njr. Mounted on /home, which is a partition on d1.
d3	Contains a directory named /net/freddie/home that is NFS-mounted on /net/freddie

Example 1:

- **Cross mount points** and **Follow NFS** are not selected.
- The backup selection list contains the following entry:
/usr/home
- In this case, NetBackup considers only the directories and files that are in the same file system as the backup selection list entry. It does not back up /home/njr or /net/freddie/home.

Example 2:

- **Cross mount points** and **Follow NFS** are selected.
- The backup selection list only contains a forward slash:
/
- In this case, NetBackup backs up all the files and directories in the tree, including those under /home/njr and /net/freddie/home.

To back up only `/usr` and individual files under `/`, leave `/` out of the list and separately list the files and directories you want to include. For example:

```
/usr  
/individual_files_under_root
```

Compression (policy attribute)

The **Compression** attribute specifies that the backups use the software compression that is possible, based on the policy type. Check the box to enable compression. By default, compression is disabled.

Note: The **Compression** attribute is independent of the VxFS compression state.

See [“About the backup and restore of compressed files on VxFS file systems”](#) on page 1127.

Compression reduces the size of a backup by reducing the size of files in the backup. In turn, the smaller backup size decreases the number of media that is required for storage. Compression also decreases the amount of data that travels over the network as well as the network load. However, compression increases the overhead computing on the client and increases backup time due to the time required to compress the files. The lower transfer rate that is associated with compression on the client reduces the ability of some tape devices to stream data. The effect of the lower transfer rate causes additional wear on those devices.

The savings in media and network resources continue to make compression desirable unless total backup time or client computing resources become a problem. If total backup time is a problem, consider multiplexing. The NetBackup multiplexing feature backs up clients in parallel, reducing the total time to back them up.

See [“Media multiplexing \(schedule attribute\)”](#) on page 756.

The degree to which a file can be compressed depends on the data type. A backup usually involves more than one type of data. Examples include stripped and unstripped binaries, ASCII, and the non-unique strings that repeat. Some data types are more favorable to compression.

Note: When compression is not used, the server may receive more data than the space that exists on the client. The discrepancy is due to client disk fragmentation and the file headers that the client adds. (To tell how much space a file occupies, run the `du` command. To tell how much free disk space is available, run the `df` command.)

[Table 20-17](#) describes the various factors to consider when you choose to use **Compression**.

Table 20-17 Considerations regarding compression

Item	Description
Data types that compress well	Programs, ASCII files, and unstripped binaries (typically 40% of the original size). Best-case compression: Files that are composed of the strings that repeat can sometimes be compressed to 1% of their original size.
Data types that do not compress well	Stripped binaries (usually 60% of original size). Worst-case compression: Files that are already compressed become slightly larger if compressed again.
Effect of file size	File size has no effect on the amount of compression. However, it takes longer to compress many small files than a single large one.
Client resources that are required	Compression requires client computer processing unit time and as much memory as the administrator configures.
Effect on client performance	Compression uses as much of the computer processing unit as available and affects other applications that require the computer processing unit. For fast CPUs, however, I/O rather than CPU speed is the limiting factor.
Files that are not compressed	NetBackup does not compress the following files: - Files that are equal to or less than 512 bytes, because that is the tar block size. - On UNIX clients, files with the following suffixes: .arc .arj .au .cpt .cpt.bin .F .F3B .gif .gz .hqx .hqx.bin .jpeg .jpg .lha .lzh .pak .iff .pit .pit.bin .scf .sea .sea.bin .sit .sit.bin .tiff .Y .zip .zom .zoo .z - On UNIX clients, if a compressed file has a unique file extension, exclude it from compression by adding it under the Client Settings (UNIX) properties.
Effect of using with storage units with SIS capabilities	If compressed data is written to a storage unit that has single-instance store (SIS) capabilities, the storage unit may not be able to use data deduplication on the compressed or the encrypted data. In data deduplication, only one instance of the file is stored. Subsequent instances of the file reference the single file.

Note: If compressed data is written to a storage unit that has deduplication capabilities, the storage unit may not be able to deduplicate the data.

Encryption (policy attribute)

The **Encryption** attribute determines whether the backup should be encrypted by the client. When the server initiates the backup, it passes on the **Encryption** policy attribute to the client in the backup request.

The client compares the **Encryption** policy attribute to the **Encryption** host properties for the client. If the encryption permissions for the client are set to REQUIRED or ALLOWED, the policy can encrypt the backups for that client.

See “[Encryption properties](#)” on page 88.

For additional encryption configuration information, see the [NetBackup Security and Encryption Guide](#).

Note: If encrypted data is written to a storage unit that has deduplication capabilities, the storage unit may not be able to deduplicate the encrypted data.

About NetBackup encryption options

NetBackup provides several methods for encrypting backups, as described in the following table.

Table 20-18 NetBackup encryption options

Option	Description
Client encryption	The NetBackup client encryption option is a software-based solution that encrypts the data on the client. The data is encrypted in transit and at rest. Each client manages its own encryption keys. To enable client encryption, select the backup policy Encryption attribute. See “Encryption (policy attribute)” on page 691.

Table 20-18 NetBackup encryption options (*continued*)

Option	Description
Tape drive encryption	With hardware-based tape drive encryption, an encrypting tape drive encrypts the data. The data is encrypted at rest only. A Key Management Service (KMS) server that is configured on the primary server manages encryption keys. It can either be NetBackup KMS (NBKMS) or external KMS. See the “Data at rest key management” chapter in the NetBackup Security and Encryption Guide. One method to manage the volumes for hardware-based tape encryption is to use a reserved prefix on the volume pool name. The storage device must have encrypting tape drives. The storage unit must specify the storage device that has the encrypting tape drives. The backup policy must specify the correct storage unit and volume pool. See “About reserved volume pool name prefixes” on page 464.
AdvancedDisk encryption	A plug-in in the NetBackup OpenStorage stack encrypts the data. The data is encrypted at rest only. A Key Management Service (KMS) server that is configured on the primary server manages encryption keys. It can either be NetBackup KMS (NBKMS) or external KMS. See the NetBackup AdvancedDisk Storage Solutions Guide.
Cloud storage encryption	A plug-in in the NetBackup OpenStorage stack encrypts the data. The data is encrypted at rest only (by default, NetBackup uses SSL for read and write operations). A Key Management Service (KMS) server that is configured on the primary server manages encryption keys. It can either be NetBackup KMS (NBKMS) or external KMS. See the NetBackup Cloud Administrator's Guide.
Media Server Deduplication Pool encryption	The MSDP deduplication plug-in encrypts the data. The data can be encrypted in transit and at rest or at rest only. The NetBackup deduplication plug-in manages the encryption keys. See the NetBackup Deduplication Guide.

Collect disaster recovery information for Bare Metal Restore (policy attribute)

The **Collect disaster recovery information for Bare Metal Restore** attribute specifies whether the BMR client agent runs on each client. If the attribute is enabled, the BMR client agent runs before each backup to save the configuration information of the client. The **Activity Monitor** displays the activity as a job separate from the backup.

Only policy types **MS-Windows** (for Windows clients) and **Standard** (for UNIX clients) support this policy attribute. This attribute is enabled by default when one of these policy types is used to create a policy on a primary server that is licensed for BMR.

For more information, see the [Bare Metal Restore Administrator's Guide](#).

Collect true image restore information (policy attribute) with and without move detection

The **Collect true image restore information** attribute specifies whether the policy collects the information necessary to perform a true image restore. A true image restore (TIR) restores the contents of a directory to reflect the contents of the directory at the time of an incremental or a full backup. Files that were deleted before the backup are not restored.

With the attribute enabled, a restore based on an incremental backup includes all files that were backed up since the last full backup. The restore also includes those files that were deleted at any time during that period.

NetBackup starts to collect the true image restore information with the next full or incremental backup for the policy. The true image restore information is collected for each client regardless of whether any files were changed.

NetBackup does not provide true image restores based on the time of a user backup or archive. However, NetBackup uses a user backup for a true image restore if the backup is more recent than the latest automatic full or incremental backup.

For true image incremental backups, enable **With move detection** to include the files that were moved, renamed, or newly installed in the directories. These files may be from a tar or a zip archive. (Depending on how the files were packaged and how they were installed, some newly installed files are not backed up by non-TIR incremental backups.

NetBackup detects changes by comparing path names and inode numbers with those from the previous full or incremental backup. If either the name or an inode number is new or changed, the file or directory is backed up. NetBackup begins to collect the information for move detection with the next full or incremental backup

for the policy. This first backup after the attribute is set always backs up all files, even if it is an incremental backup.

Note: With **move detection** must be enabled to create a synthetic backup.

See [“Synthetic backup \(schedule attribute\)”](#) on page 741.

The following examples show how move detection backs up the files that otherwise would not be backed up:

- On Windows:
 - A file that is named `C:\pub\doc` is moved to or installed in `C:\spec\doc`. The archive bit is unchanged but `C:\spec\doc` is new in the `C:\spec\` directory and is backed up.
 - A directory that is named `C:\security\dev\` is renamed as `C:\security\devices\`. The archive bit is unchanged but `C:\security\devices\` is a new directory and is backed up.
- On UNIX:
 - A file that is named `/home/pub/doc` is moved to `/home/spec/doc`. The modification time is unchanged but `/home/spec/doc` is new in the `/home/spec/` directory and is backed up.
 - A directory that is named `/etc/security/dev` is renamed as `/etc/security/devices`. The modification time is unchanged but `/etc/security/devices` is a new directory and is backed up.
 - A file that is named `/home/pub/doc` is installed when extracted from a UNIX `tar` file. The modification time is before the time of the last backup. The file is new in the `/home/pub/` directory and is backed up.
 - A file that is named `docA` is removed, and then a file that is named `docB` is renamed as `docA`. The new `docA` has the same name but since its inode number has changed, it is backed up.

NetBackup begins to collect the information that is required for move detection with the next full or incremental backup for the policy. This first backup after the attribute is set always backs up all files, even if it is an incremental backup.

Move detection consumes space on the client and the backup can fail if there is not enough disk space available.

Example of true image restores

The following table lists the files that were backed up in the `/home/abc/doc/` directory of a UNIX client during a series of backups between 12/01/2015 and 12/04/2015. **Collect true image restore information** was enabled for the policy that performed the backups.

Table 20-19 Sample backups taken before a true image restore

Day	Type of backup	Files that are backed up in <code>/home/abc/doc</code>
12/01/2015	Full	file1 file2 dirA/fileA dirB/fileB file3
12/02/2015	Incremental	file1 file2 dirA/fileA -----
12/03/2015	Incremental	file1 file2 dirA/fileA -----
12/04/2015	User backup	file1 file2 dirA/fileA ----- dirC/fileC file4
12/04/2015	Incremental	file1 file2 ----- file4

Note: Dashes (-----) indicate that the file was deleted before this backup.

A restore of the 12/04/2015 version of the `/home/abc/doc/` directory produces following results:

After a regular restore The restored directory contains all files and directories that ever existed in `/home/abc/doc/` from 12/01/2015 (last full backup) through 12/04/2015:

```
file1
file2
dirA/fileA
dirB/fileB
file3
dirC/fileC
file4
```

After a true image restore

The restored directory contains only the files and directories that existed at the time of the incremental backup:

```
file1
file2
file4
```

NetBackup does not restore any of the files that were deleted before the 12/04/2015 incremental backup.

The restored directory does not include the subdirectories `dirA` and `dirC`, even though they were backed up on 12/04/2015 with a user backup.

NetBackup did not restore these directories because they did not exist at the time of the incremental backup. The incremental backup was the reference for the true image restore.

Consider the following points to use either **Collect true image restore** or **Collect true image restore with move detection**:

- NetBackup collects additional information for the incremental backups that collect true image restore information. Policies that use move detection require even more space.
- Incremental backups are slower for a policy in which true image restore information is collected.
- Configure the period of time that NetBackup retains the true image restore information. Set the **Keep true image restoration (TIR) information** property in the **Clean-up** properties dialog box.
See [“Clean up properties”](#) on page 61.
- Only directories can be listed and selected. In true image restore mode, the client interface does not display individual files. Refer to the online Help in the **Backup, Archive, and Restore** client interface for more information on true image restores.
- A true image restore preserves the files that are currently in the directory but were not present when the backup was completed. If you created a file `file5` after an incremental backup on 12/04/2015 but before a restore, the contents of the restored directory would be as follows:

```
file1
file2
file4
file5
```

Allow multiple data streams (policy attribute)

The **Allow multiple data streams** attribute specifies that NetBackup can divide automatic backups for each client into multiple jobs. The directives, scripts, or templates in the backup selection list specify whether each job can back up only a part of the backup selection list. Because the jobs are in separate data streams, they can occur concurrently.

The directives, scripts, or templates in the backup selection list determine the number of streams (backup jobs) that start for each client. The list also determines how the backup selection list is divided into separate streams.

The following settings determine the number of streams that can run concurrently:

- Number of available storage units
- Multiplexing settings
- Maximum jobs parameters

Multistreamed jobs consist of a parent job to perform stream discovery and children jobs for each stream. Each child job displays its own job ID in the **Job ID** column in the **Activity Monitor**. The job ID of the parent job appears in the **Parent Job ID** column, which is not displayed by default. Parent jobs display a dash (-) in the **Schedule** column.


Note: If this attribute is enabled, and a file system is in a client's exclude list, a NetBackup job appears in the **Activity Monitor** for the excluded file system. However, no files in the excluded file system are backed up by the job.

The following table describes the reasons to use multiple data streams.

Table 20-20 Reasons to use multiple data streams

Reason	Description
To reduce backup time	Multiple data streams can reduce the backup time for large backups by splitting the backup into multiple streams. Use multiplexing, multiple drives, or a combination of the two to process the streams concurrently. Configure the backup so that each device on the client is backed up by a separate data stream that runs concurrently with streams from other devices. For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. The heads must move back and forth between the tracks that contain files for the respective streams. Figure 20-5 shows why multiple concurrent streams from a single device are not recommended.
To reduce retry time for backup failures	Because the backup streams run independently, the use of multiple data streams can shorten the retry time in the event of a backup failure. A single failure only terminates a single stream. NetBackup can restart the failed stream without restarting the others. For example, assume the backup for a 10-gigabyte partition is split into five streams, each containing 2 gigabytes. If the last stream fails after it writes 1.9 gigabytes (a total of 9.9 gigabytes is backed up), NetBackup retries only the last gigabyte stream. If the 10-gigabyte partition is backed up without multiple data streams and a failure occurs, the entire 10-gigabyte backup must be retried. The Schedule backup attempts property in the Global attributes properties applies to each stream. For example, if the Schedule backup attempts property is set to 3, NetBackup retries each stream a maximum of three times. The Activity monitor displays each stream as a separate job. Use the job details view to determine the files that are backed up by each of these jobs. See “Global attributes properties” on page 109.
To reduce administration by running more backups with fewer policies	Use multiple data streams in a configuration that contains large file servers with many file systems and volumes. Multiple data streams provide more backups with fewer policies than are otherwise required.

Figure 20-5 Multiple stream recommendations



The following table describes the aspects of multiple data streams that are adjustable.

Table 20-21 Adjustable aspects of multiple data streams

Item	Description
The total number of streams	The backup selection list determines the total number of streams that are started. The NEW_STREAM directive lets you configure a fixed number of streams, or you can allow the client dynamically define the streams. See “About the directives on the Backup Selections list” on page 810. Note: For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. Backup times are affected because the device heads must move between the tracks that contain files for the respective streams.

Table 20-21 Adjustable aspects of multiple data streams (*continued*)

Item	Description
The number of streams that run concurrently	The following factors determine the number of streams that can run concurrently for a policy or client: ■ Number of the drives that are available. ■ Maximum concurrent jobs settings for the policy and client. ■ Storage unit and schedule multiplexing limit. Each storage unit and each schedule have a maximum multiplexing setting. The lower of the two settings is the limit for a specific schedule and storage unit. The maximum streams are limited to the sum of the multiplexing limits for all drives available in the storage unit and schedule combinations. For example, assume that two storage units have one drive in each. Multiplexing on storage unit 1 is set to 3 and multiplexing on storage unit 2 is set to 5. If multiplexing is set to 5 or greater in the schedules, 8 streams can run concurrently. See “Media multiplexing (schedule attribute)” on page 756.

The maximum jobs settings limit the maximum number of streams as follows:

Table 20-22 Job settings that limit the maximum number of streams

Item	Access method
Maximum jobs per client (host property)	■ On the left select Hosts > Host properties. ■ Select the primary server and select Edit primary server. ■ Select Global attributes. See “Global attributes properties” on page 109. See “Media multiplexing (schedule attribute)” on page 756.
Limit jobs per policy (policy attribute)	■ In the left pane, expand Protection > Policies. ■ Open the policy that you want to modify. See “Limit jobs per policy (policy attribute)” on page 678.
Maximum data streams (host property)	■ On the left select Hosts > Host properties. ■ Select the primary server and select Edit primary server. ■ Select Client attributes. See “General tab of the Client attributes properties” on page 66.

Job settings also affect the maximum number of streams. The following table describes the interdependency of these settings.

Table 20-23 Interdependency of job settings

Item	Description
Maximum data streams property is disabled.	NetBackup uses the value for either Maximum jobs per client or Limit jobs per policy , whichever is lower.
Maximum data streams property is enabled.	NetBackup ignores Maximum jobs per client . Instead, NetBackup uses the value for either Maximum data streams or Limit jobs per policy , whichever is lower.

See [“About the directives on the Backup Selections list”](#) on page 810.

Use multiple MSDP nodes (policy attribute)

Available in NetBackup 11.0 and later, the policy attribute **Use multiple MSDP nodes** improves load balancing for multistream backups on MSDP cluster storage servers. This feature enables the distribution of backup streams across the multiple MSDP nodes, optimizing throughput, and reducing potential congestion on individual nodes.

The **Use multiple MSDP nodes** attribute is available on the NetBackup web UI for the following policy types:

- Standard
- MS-Windows
- MS-SQL-Server

The policy attribute is not available for the following policy types:

- Oracle
- Epic-Large-File
- Universal-Share
- MSDP-Object-Store

For more information, see the *NetBackup Deduplication Guide*.

Client-side deduplication (policy attribute)

The **Client-side deduplication** attribute appears only if the NetBackup Data Protection Optimization Option license is active.

The following options are available when you configure policies:

Use individual client settings configured in host properties	The client uses Deduplication setting that is configured for it in the host properties.
Disable for all clients	The clients do not deduplicate their own data and do not send their backup data directly to the storage server. The NetBackup clients send their data to a deduplication media server. That server deduplicates the data and then sends it to the storage server.
Enable for all clients	The clients deduplicate their own data. They also send it directly to the storage server. Media server deduplication and data transport are bypassed.

The **Deduplication** property configures clients for client-side deduplication. The **Client-side deduplication** attribute overrides the **Deduplication** property. The **Deduplication** property is found on the **General** tab of the **Client attributes** host properties.

See [“Where deduplication should occur”](#) on page 69.

See the [NetBackup Deduplication Guide](#).

Enable granular recovery (policy attribute)

The **Enable granular recovery** attribute is available for the following policy types:

- MS-Exchange-Server
- MS-SharePoint
- MS-Windows (for Active Directory)

With this option enabled, users can restore the individual objects that reside within a database backup image, such as:

- A user account from an Active Directory database backup
- Email messages or folders from an Exchange database backup
- A document from a SharePoint database backup

NetBackup does not support the compression or encryption of backups that use Granular Recovery Technology (GRT). When the **Enable granular recovery** option is enabled, the **Compression** option and the **Encryption** option are automatically disabled.

Granular-level restores can be performed only if the backup was written to a disk storage unit.

For more information on how to configure NetBackup to perform granular-level backups, see the following:

- [NetBackup for SharePoint Server Administrator's Guide](#)
- [NetBackup for Exchange Server Administrator's Guide](#)
- See "[Active Directory granular backups and recovery](#)" on page 835.

Use Accelerator (policy attribute)

NetBackup Accelerator increases the speed of full backups. The increase in speed is made possible by change detection techniques on the client. The client uses the change detection techniques and the client's current file system to identify the changes that occurred since the last backup. The client sends the changed data to the media server in a more efficient backup stream. The media server combines the changed data with the rest of the client's data that is stored in previous backups.

If a file or portion of a file is already in storage and has not been changed, the media server uses the copy in storage rather than reading it from the client. The end result is a full NetBackup backup.

Accelerator has the following advantages:

- Reduces the I/O and CPU overhead on the client. The result is a faster backup and less load on the client.
- Creates a compact backup stream that uses less network bandwidth between client and server.
- Creates a full image that contains all data that is needed for restore.

Note: Accelerator operates differently when used for backup of virtual machines, NDMP, and databases.

For full details on Accelerator for VMware, see the [NetBackup for VMware Administrator's Guide](#).

For full details on Accelerator for NDMP, see the [NetBackup for NDMP Administrator's Guide](#).

For full details on Accelerator for Oracle, SharePoint, Exchange, or SQL Server, see the [NetBackup guide for that agent](#).


How the NetBackup Accelerator works

The NetBackup Accelerator creates the backup stream and backup image as follows:

- If the client has no previous backup, NetBackup performs a full backup and creates a track log. The track log contains information about the client's data, for comparison at the next backup.
- At the next backup, NetBackup identifies data that has changed since the previous backup. To do so, it compares information from the track log against information from the file system for each file. For NTFS and ReFS file systems, it also uses the Windows change journal to help identify the data that has changed since the last backup.
Accelerator uses the Windows change journal in two ways: To check for changes in the file system metadata, and to help detect which files have changed since the last backup.
See [“Accelerator and the Windows change journal”](#) on page 706.
- The NetBackup client sends to the media server a backup stream that consists of the following: The client's changed blocks, and the previous backup ID and data extents (block offset and size) of the unchanged blocks.
- The media server receives the client's changed blocks and the backup ID and data extents of the unchanged blocks. From the backup ID and file system descriptors, the media server locates the rest of the client's data in existing backups.
- The media server directs the storage server to write the changed blocks and the unchanged blocks in a new full image.

[Figure 20-6](#) shows how an Accelerator backup stream is composed.

Figure 20-6 NetBackup client: Accelerator backup stream



[Figure 20-7](#) shows how the media server constructs a backup image from the Accelerator stream and from previous backups:

Figure 20-7 NetBackup media server constructs backup image





Figure 20-8 recaps Accelerator details in the context of the principal NetBackup processes.

Figure 20-8 Process overview of Accelerator backup



Accelerator and the Windows change journal

For Windows NTFS and ReFS file systems, the Accelerator uses the Windows change journal to help identify the files that changed since the previous backup.

The extent to which the Windows change journal is used depends on the following:

- Whether the **Use Change Journal** option has been enabled for the NetBackup client.
- Whether the policy contains a schedule with the **Accelerator forced rescan** option enabled.

Note: Regardless of the **Use Change Journal** setting or the **Accelerator forced rescan** setting, the Windows change journal always helps determine if a file has changed.

Table 20-24 How Accelerator uses the Windows change journal

NetBackup configuration	Accelerator use of change journal
The Use Change Journal option is not enabled	For full backups and incremental backups, NetBackup examines the metadata for every file to determine if the file has changed. Even without the Use Change Journal option, the Windows change journal helps determine if a file has changed.
The Use Change Journal option is enabled but the Accelerator forced rescan option is not enabled on any of the policy's schedules	The Windows change journal helps determine if a file has changed during both full backups and incremental backups. ■ During incremental backups, the Windows change journal is used exclusively to determine if a file has changed. NetBackup does not examine the file metadata. ■ During full backups, NetBackup examines the metadata for every file, to determine whether the file has changed. In addition, NetBackup uses the Windows change journal to help determine whether a file has changed.
The Use Change Journal option is enabled and the Accelerator forced rescan option is enabled on at least one of the policy's schedules	The Windows change journal helps determine if a file has changed during both full backups and incremental backups. ■ During incremental backups, the Windows change journal is used exclusively to determine if a file has changed. NetBackup does not examine the file metadata. ■ During full backups: ■ For a full backup schedule that does not enable the Accelerator forced rescan option: The Windows change journal is used exclusively to determine if a file has changed. NetBackup does not examine the file metadata. ■ For a full backup schedule that enables the Accelerator forced rescan option: NetBackup examines the metadata for every file, to determine whether the file has changed. In addition, NetBackup uses the Windows change journal to help determine whether a file has changed.

Accelerator notes and requirements

Note the following about the NetBackup Accelerator:

- NetBackup Accelerator requires the Data Protection Optimization Option license. For the latest information on licensing, contact your NetBackup sales or partner representative.

- Supports the disk storage units only. Supported storage includes **Media Server Deduplication Pool**, NetBackup appliance, cloud storage, and qualified third-party OST storage.

For supported storage types, see the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List* at the following URL:

https://www.veritas.com/support/en_US/article.100040093

- Storage unit groups are supported only if the storage unit selection in the group is Failover.
- Supports the MS-Windows and Standard policy types. Supports all features of NetBackup that work with the MS-Windows or Standard policy types.

Note: Accelerator operates differently when used for backup of virtual machines, NDMP, or Oracle databases. For full details, see the following:

NetBackup for VMware Administrator's Guide

NetBackup for NDMP Administrator's Guide

NetBackup for Oracle Administrator's Guide

- Accelerator backups ignore the policy attribute that resets the `atime` on files after backup.
By default, NetBackup records the file access time (`atime`) for each UNIX file that it backs up, and then resets the `atime` after the file is backed up. Resetting the `atime` in this way causes the change time (`ctime`) to update as well. The **Reset file access time to the value before backup** policy attribute normally controls this behavior (the default is enabled).
When NetBackup Accelerator is used to perform backups, NetBackup does not reset the `atime` after the file is backed up, regardless of the policy attribute setting. NetBackup does not reset the `atime` (which avoids updating the `ctime`) because a `ctime` update would require a loss in Accelerator optimization.
See ["Client settings properties for UNIX clients"](#) on page 73.
- Supports the full backups and incremental backups.
See ["Accelerator backups and the NetBackup catalog"](#) on page 710.
- Supports all platforms, file systems, and logical volumes that NetBackup supports.
- Supports the Windows NTFS and ReFS change journal (**Use Change Journal**) but does not support the VxFS change journal.
- For every policy that enables the **Use Accelerator** option, the following backup schedules are recommended at a minimum:

A full backup schedule with the **Accelerator forced rescan** option enabled.
 Another full backup schedule without the **Accelerator forced rescan** option enabled.

See [“Accelerator forced rescan option \(schedule attribute\)”](#) on page 742.

- If **Collect true image restore information** is enabled in the policy, NetBackup performs a regular file system scan to determine the files and folders to include in the backup. It also queries the change journal to help determine which files have changed.

More information is available on the **Use change journal** option and the Accelerator:

See [“Accelerator and the Windows change journal”](#) on page 706.

- If a previous backup of the client does not exist, NetBackup performs a full backup and creates a track log on the client. This initial backup occurs at the speed of a normal (not accelerated) full backup. Subsequent Accelerator backups of the client use the track log for accelerated backup speed.

Note: When you first enable a policy to use Accelerator, the next backup (whether full or incremental) is in effect a full backup: It backs up all files in the **Backup Selections** tab. If that backup was scheduled as an incremental, it may not complete within the backup window.

- NetBackup retains track logs for future Accelerator backups. Whenever you change the policy's backup selections list, NetBackup does a full non-accelerated backup of the selections that were changed or added in the list. The unchanged backup selections are processed as normal Accelerator backups.
- If the storage unit that is associated with the policy cannot be validated when you create the policy, note: The storage unit is validated later when the backup job begins. If Accelerator does not support the storage unit, the backup fails. In the `bpbrrm` log, a message appears that is similar to one of the following:

```
Storage server %s, type %s, doesn't support image include.
```

```
Storage server type %s, doesn't support accelerator backup.
```

- Accelerator requires that the storage have the `OptimizedImage` attribute enabled.
- The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.

Accelerator backups and the NetBackup catalog

Use of Accelerator does not affect the size of the NetBackup catalog. A full backup with Accelerator generates the same catalog size as a full backup of the same data without Accelerator. The same is true of incremental backups: use of Accelerator does not require more catalog space than the same backup without Accelerator.

However, if you enable **Calculate file hash** option on the web UI, the NetBackup catalog is expected to increase by 20% or more. This option is used to save the file hash information to the NetBackup catalog. See the *NetBackup Web UI Administrator's Guide* for more information.

A potential catalog effect does exist, depending on how often you use Accelerator with full backups. A full backup with Accelerator completes faster than a normal full. It may therefore be tempting to replace your incremental backups with Accelerator full backups. Note: Since a full backup requires more catalog space than an incremental, replacing incrementals with fulls increases the catalog size. When changing your incrementals to fulls, you must weigh the advantage of Accelerator fulls against the greater catalog space that fulls require compared to incrementals.

See [“Configuring Accelerator”](#) on page 710.

Configuring Accelerator

The following table outlines the procedure to configure the full backups that use the NetBackup Accelerator.

Table 20-25 To configure Accelerator on full backups

Task	Procedure and notes
Make sure that you have a storage unit that supports Accelerator.	Accelerator supports disk storage units only. Supported storage unit types are PureDisk (Media Server Deduplication Pool and NetBackup appliance), cloud storage, and qualified third-party OST storage. The NetBackup device mapping files contain a complete list of supported storage units (see the next task). Note: Storage unit groups are supported only if the storage unit selection in the group is Failover.
Update the NetBackup device mapping files if needed.	The NetBackup device mapping files contain all storage device types that NetBackup can use. To add support for the new devices or upgraded devices that support Accelerator, download the current device mapping files from the Cohesity Technical Support. See “About the device mapping files” on page 419. See “Downloading the device mapping files” on page 419.

Table 20-25 To configure Accelerator on full backups (*continued*)

Task	Procedure and notes
Configure a Standard , MS-Windows , VMware , NDMP , or Oracle backup policy.	Select the following on the policy Attributes tab: ■ A disk storage unit that supports Accelerator. ■ The Use Accelerator option. Accelerator operates differently when used for backup of virtual machines, NDMP, or Oracle databases. For full details, see the following guides: ■ The <i>NetBackup for VMware Administrator's Guide</i>. ■ The <i>NetBackup for NDMP Administrator's Guide</i>. ■ The <i>NetBackup for Oracle Administrator's Guide</i>. (Optional) To save the file hash to the NetBackup catalog, enable the option Calculate file hash. This option is only available when you configure a policy in the web UI. It is only available for the policy types Windows, Standard, and NAS-Data-Protection. See the <i>NetBackup Web UI Administrator's Guide</i> for configuration instructions.
To periodically establish a new baseline of change detection on the client, you can select the Accelerator forced rescan option on the Schedule Attribute tab of the policy.	This option provides an additional level of Accelerator change detection in the client's data. This option reduces the speed of Accelerator. See "Accelerator forced rescan option (schedule attribute)" on page 742. See "Accelerator and the Windows change journal" on page 706.

Accelerator messages in the backup job details log

A NetBackup backup that uses Accelerator writes a message similar to the following in the job details log:

```
11/23/2011 10:50:27 AM - Info bpbmr(pid=412) accelerator enabled
```

When NetBackup uses the NTFS or ReFS change journal for the backup, messages similar to the following appear in the job details log:

```
9/24/2011 8:54:14 PM - Info bpbkar32(pid=7868) change journal enabled
for <C:\>
```

```
9/24/2011 8:54:14 PM - Info bpbkar32(pid=7868) using change journal
data for <C:\>
```

If the NTFS or ReFS change journal cannot be used, a message of the following form appears in the job details log:

```
not using change journal data for <backup selection>: <reason ...>
```

See [“Log messages about the Use Change Journal option and Accelerator”](#) on page 712.

When the **Accelerator forced rescan** option is used, a message similar to the following appears in the job details log:

```
9/25/2011 5:46:52 PM - Info bpbrm(pid=4136) Accelerator enabled
backup with checksum based change detection needs to read each file and
calculate the checksum, and will have longer backup time.
```

See [“Log messages about the Use Change Journal option and Accelerator”](#) on page 712.

NetBackup logs for Accelerator

For log messages about Accelerator, see the following NetBackup log directories.

Table 20-26 NetBackup logs that pertain to Accelerator

Log directory	Resides on
Windows: <i>install_path</i> \NetBackup\logs\bpbrm UNIX: /usr/opensv/netbackup/logs/bpbrm	NetBackup primary or media server
Windows: <i>install_path</i> \NetBackup\logs\bptm UNIX: /usr/opensv/netbackup/logs/bptm	NetBackup media server
Windows: <i>install_path</i> \NetBackup\logs\bpbkar UNIX: /usr/opensv/netbackup/logs/bpbkar	NetBackup client

To create the log directories, run the following command on the NetBackup servers and client:

On Windows:

```
install_path\NetBackup\logs\mklogdir.bat
```

On UNIX:

```
/opt/opensv/netbackup/logs/mklogdir
```

Log messages about the Use Change Journal option and Accelerator

The Accelerator uses the NTFS or ReFS change journal to help identify data changes on the Windows client. The following table describes the change journal

messages that may appear in the NetBackup job details log in the Activity Monitor. The left column lists the messages.

For Accelerator backups, these messages indicate various impediments to the use of the Windows change journal. In general, they indicate either of the following:

- That the Accelerator cannot use the change journal for the backup. To use the change journal, you may need to reconfigure the policy.
- That the Accelerator cannot exclusively use the change journal to detect changes in the file system. NetBackup examines the file system metadata for each file, to detect which files have changed. It also uses the change journal to help detect which files have changed.

Note: In the following messages, the variable <%%s> represents the items in your backup selections list.

The change journal messages often appear in pairs. The first message indicates why the change journal cannot be used. For example:

```
not using change journal data for <%%s>: forced rescan requested
```

The second message indicates that the Accelerator nevertheless can use the change journal to detect changed files:

```
not using change journal data for enumeration for <%%s> but will use  
it for change detection
```

Table 20-27 Accelerator messages on the Windows **Use Change Journal** option (job details log)

Message in NetBackup job details	Description
not using change journal data for enumeration for <%%s> but will use it for change detection	NetBackup uses the change journal to help detect which files have changed. NetBackup also examines the file system metadata for each file to help detect which files have changed. Note: Before this message appears, another message explains why NetBackup does not rely entirely on the change journal data to detect changed files.
not using change journal data for <%%s>: forced rescan requested	The Accelerator forced rescan option is enabled on the full backup schedule for the policy. As a result, NetBackup cannot rely entirely on the change journal data to detect changed files. NetBackup also examines the file system metadata for each file to help detect which files have changed.

Table 20-27 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: filter checksum calculation failed	The backup uses a number of filters to determine which files and directories to include in the backup. The filters are the following: NetBackup exclude and include lists, the files and directories that are included in the Shadow Copy Components and in the system state backup, and others. During a backup, a checksum is calculated against the filters. When a new backup runs, the checksum for the current backup is compared against the checksum of the previous backup. If the checksums do not match, the change journal data cannot be used. Instead, Accelerator performs a regular file system scan to determine the files and folders to include in the backup. No action is required. If the filters do not change between backups, the change journal data is used at the next backup.
not using change journal data for <%%s>: unable to validate change journal usage <reason=%%s>	Accelerator cannot use the Use Change Journal option in the following cases: ■ No previous backup exists: No baseline update sequence number (USN) was established from the NTFS change journal. ■ The previous backup was not successful. In these cases, the Use Change Journal option cannot be used. Accelerator performs a regular file system scan to determine the files and folders to include in the backup. No action is required. Accelerator uses the Use Change Journal option at the next backup if all conditions are met.
not using change journal data for <%%s>: unable to initialize change journal usage <reason=%%s>	Accelerator cannot use the Use Change Journal option in the following cases: ■ Too much activity has occurred on the system (records were purged from the change journal databases before they could be processed). ■ Data corruption occurred. Instead, Accelerator performs a regular file system scan to determine the files and folders to include in the backup. No action is required. When records have been purged, a new baseline is created when the current backup completes. If corruption existed, it is detected and the databases are re-created. Accelerator uses the Use Change Journal option at the next backup.
not using change journal data for <%%s>: hard link or reparse point change detected	Accelerator cannot use the Use Change Journal option if a change corresponds to a hard link or to a reparse point. Note that the change may correspond to any number of files and directories and the NTFS change journal does not track all of them. No action is required. If no further changes occur to hard links or reparse points, the Use Change Journal option can be used at the next backup.

Table 20-27 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: not supported with true image backups	Collect true image restore information or Collect true image restore information with move detection is specified in the policy. To process the files and determine which files have changed, NetBackup performs a regular file system scan to determine the files and folders to include in the backup. It also queries the change journal to help determine which files have changed.
not using change journal data for <%%s>: invalid schedule type	Accelerator does not support the selected schedule type with the Use Change Journal option. The Use Change Journal option is supported for incremental backups (cumulative or differential) or full backups. For full backups, Use Accelerator must be enabled on the policy Attributes tab.
not using change journal data for <%%s>: path must be local and not contain directory junctions and/or symbolic links	In the backup selections, a path contains a reparse point (directory junction or a symbolic link). The Use Change Journal option cannot be used. NetBackup must perform a regular file system scan to back up the directories correctly.
not using change journal data for <%%s>: change journal setup is not complete (may still be in progress)	The Use Change Journal option was recently enabled. After Use Change Journal is first enabled, the NetBackup client may need time to process the NTFS change journal and set up its databases. The Use Change Journal option may be ready at the next backup.
not using change journal data for <%%s>: unable to locate journal data	The Use Change Journal option was recently enabled. After Use Change Journal is first enabled, the NetBackup client may need time to process the NTFS change journal and set up its databases. The Use Change Journal option may be ready at the next backup.
not using change journal data for <%%s>: invalid change journal data	When many changes occur on a volume, the NetBackup Client Service may need to increase the size of the internal databases. As a result, the databases may become invalid. After the databases are increased in size and are synchronized with the NTFS change journal, they are marked as valid. The Use Change Journal option is used with the next backup.

Table 20-27 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: unable to apply snapshot	For a snapshot-based backup, NetBackup uses the change journal databases on the snapshot instead of on the live volume. An error occurred when NetBackup attempted to open the databases on the snapshot. When the next backup runs, NetBackup creates a new snapshot and the databases may be opened without error.
not using change journal data for <%%s>: no previous track log	No previous full backup exists that used Accelerator. NetBackup supports the Use Change Journal option with Accelerator on a full backup only if a previous full backup exists that used Accelerator.
not using change journal data for <%%s>: not supported with regular full backups	The Use Accelerator option is not enabled on the policy. For full backups, the Use Change Journal option requires that Use Accelerator is enabled.
not using change journal data for <%%s>: unable to validate change journal usage <reason=previous backup wasn't a successful backup>	If a backup is partially successful (status code 1), the next Accelerator backup cannot use the Use Change Journal option. NetBackup can use the Use Change Journal option after the next successful backup.
not using change journal data for <%%s>: not supported	A backup selection in the policy is a resource for which the Windows change journal is not supported. Examples of unsupported resources are FAT volumes and FAT32 volumes.
not using change journal data for <%%s>: not supported for remote frozen images	The backup has attempted to use a remote frozen image. The Use Change Journal option is not supported with remote frozen images.
not using change journal data for <%%s>: not enabled	The Use Change Journal option is not enabled.

Table 20-27 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: not configured for use	The Use Change Journal option is not enabled.
not using change journal data for <%%s>: unable to determine volume guid	An internal error occurred. The volumes to back up should be checked to ensure that a volume GUID is associated with each volume.
not using change journal data for <%%s>: snapshot has not been applied (unable to track open files)	To use the change journal data, NetBackup must be able to take a snapshot of the backup volume to correctly identify and handle open files. This error occurs if a backup runs before NetBackup can take a snapshot of the volume. If a snapshot can be taken before the next backup occurs, NetBackup may be able to use change journal data at the next backup.
not using change journal data for <%%s>: invalid policy type	The Use Change Journal option is only supported for Standard or MS-Windows policies.
not using change journal data for <%%s>: list of accelerator based backups does not match list of change journal based backups	On the client, NetBackup keeps track of both Accelerator backups and change journal-based backups. If the lists of these backups do not match, one of the following occurred: ■ An Accelerator backup occurred that did not use the change journal. ■ A change journal-based backup occurred that did not use the Accelerator. If either case, the Use Change Journal option cannot be used until the next full backup occurs.
not using change journal data for <%%s>: previous backup had change journal data that was not complete (missing usn records)	The change journal databases are fixed in size: they can contain only a fixed number of records. If the volume undergoes a lot of activity, records may be purged from the change journal database before a backup can process them. In that case, the Use Change Journal option cannot be used. No action is required. Accelerator uses the Use Change Journal option at the next backup if all conditions are met.

Table 20-27 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: not supported for non-local volumes / file systems	The backup selection in the policy is not for a local volume. For example, the backup selection specifies a UNC path.
not using change journal data for <%%s>: no forced rescan schedule configured	This message appears only during Accelerator full backups. To use the change journal data during Accelerator full backups, a policy schedule with the Accelerator forced rescan option is required.
not using change journal data for <%%s>: forcing rescan, each file will be read in order to validate checksums	This message appears only during full backups, when the policy is not configured with the Accelerator forced rescan option. Before this message appears, another message explains why the change journal data cannot be used.

About reporting the amount of Accelerator backup data that was transferred over the network

For Accelerator backup reporting, several NetBackup commands can report the amount of data that is transferred over the network for each Accelerator backup. The amount of transferred data is often much less than the size of the Accelerator backup image.

For each Accelerator backup, NetBackup combines the client's (or VM's) changed blocks with the unchanged data from previous backups to synthesize a backup image. However, NetBackup sends only the changed data over the network when the backup occurs. The resulting backup image may be much larger than the amount of backup data that travels the network. For backup reporting, it may be important to distinguish between the backup image size and the amount of data that was transferred over the network.

For Accelerator backups, the network-transferred data can appear in the output of the following NetBackup commands: `bpdjobs`, `bpimagelist`, and `bpclimagelist`.

[Table 20-28](#) lists the default location of these commands.

Table 20-28 Default location of `bpdbjobs`, `bpimagelist`, and `bpclimagelist`

Command	Default location
<code>bpdbjobs</code> , <code>bpimagelist</code>	Windows: <i>install_path\NetBackup\bin\admincmd\</i> UNIX, Linux <i>/usr/opensv/netbackup/bin/admincmd/</i>
<code>bpclimagelist</code>	Windows: <i>install_path\NetBackup\bin\</i> UNIX, Linux <i>/usr/opensv/netbackup/bin/</i>

The following example uses the `bpimagelist` command to show the results of a backup of `acmevm2`:

```
bpimagelist -backupid acmevm2
```

Example output:

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 7799632 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

In this example, the backup image size in kilobytes is 7799632, and the amount of data that was transferred over the network is 225792.

You can use the following commands to show the amount of data that was transferred over the network for an Accelerator backup.

bpimagelist

```
bpimagelist -backupid backup_id [-l | -L | -json | -json_compact]
```

Brackets [] indicate optional elements, and the vertical bars | indicate that you can choose only one of the options within the brackets.

[Table 20-29](#) describes how the network-transferred data field appears in the `bpimagelist` output.

Table 20-29 The **bpimagelist** options that show the amount of network-transferred data for Accelerator backups

bpimagelist option	How the network-transferred data field appears
No option	The field is unlabeled. For example: 225792 See the <code>bpimagelist</code> example output earlier in this topic.
-l	The field is unlabeled (same as no option). For example: 225792
-L	The field is labeled. For example: Kilobytes Data Transferred: 225792
-json	The field is labeled. For example: "kilobytes_data_transferred": 225792,
-json_compact	The field is labeled. For example: "kilobytes_data_transferred":225792,

bpdbjobs

```
bpdbjobs -jobid job_id -report -most_columns
```

or

```
bpdbjobs -jobid job_id -report -all_columns
```

The network-transferred data field appears at the end of the output.

bpclimagelist

```
bpclimagelist -client client_name
```

This command can only show the network-transferred data in the field that normally shows the Accelerator backup image size. To show the network-transferred data with this command, you must configure a NetBackup setting:

See [“Replacing the Accelerator image size with the network-transferred data in NetBackup command output”](#) on page 721.

Additional details on these commands are available in the *NetBackup Commands Reference Guide*.

Replacing the Accelerator image size with the network-transferred data in NetBackup command output

You can configure the output of `bpimagelist`, `bpdbjobs`, and `bpclimagelist` to show the amount of Accelerator backup data that was transferred over the network instead of the backup image size.

The following is the default `bpimagelist` output that shows the Accelerator image size (see the circled value 7799632). The amount of network-transferred data appears farther down in the output (225792):

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 7799632 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

You can configure NetBackup command output to show the network-transferred data in the image size field. In the output, the image size value is replaced with the network-transferred data value (see the following example). A script that reads the image size from the command output now reads the amount of network-transferred data.

In the following `bpimagelist` output, the image size field shows the network-transferred data (225792):

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 225792 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

Note: The same change occurs in the labeled output of the commands (such as with the `-L` option of `bpimagelist`). For example, the `Kilobytes` field shows the transferred data value (225792 in the example) rather than the Accelerator backup image size.

To enable the reporting of network-transferred data in the Accelerator image size field of `bpimagelist`, `bpdbjobs`, and `bpclimagelist`

- ◆ Use the `bpsetconfig` command to enable the output change.

To enable this change for the `bpclimagelist` command, enter the `bpsetconfig` command on the primary server. To enable this change for `bpimagelist` or `bpdbjobs`, enter the `bpsetconfig` command on the server where you intend to run `bpimagelist` or `bpdbjobs`.

Refer to [Table 20-30](#) for the `bpsetconfig` command to use based on the type of Accelerator backup that you want to report on.

Table 20-30 To enable the reporting of network-transferred data in the Accelerator image size field of `bpimagelist`, `bpclimagelist`, or `bpdbjobs` output

Type of backup to report on	Enter this command
Incremental Accelerator virtual machine backups (VMware and Hyper-V)	Windows <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_VIRTUAL install_path\NetBackup\bin\admincmd\bpsetconfig</pre> UNIX, Linux <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_VIRTUAL" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>
All Accelerator virtual machine backups (VMware and Hyper-V, full and incremental)	Windows <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_VIRTUAL install_path\NetBackup\bin\admincmd\bpsetconfig</pre> UNIX, Linux <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_VIRTUAL" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>

Table 20-30 To enable the reporting of network-transferred data in the Accelerator image size field of **bpimagelist**, **bpclimagelist**, or **bpdbjobs** output (*continued*)

Type of backup to report on	Enter this command
All incremental Accelerator backups (physical clients and virtual machines)	Windows <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_ALL install_path\NetBackup\bin\admincmd\bpsetconfig</pre> UNIX, Linux <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_ALL" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>
All Accelerator backups (full and incremental, physical clients and virtual machines)	Windows <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_ALL install_path\NetBackup\bin\admincmd\bpsetconfig</pre> UNIX, Linux <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_ALL" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>

To reset the command output to the default setting

- ◆ To disable the reporting of network-transferred data in the Accelerator image size field (return to default), enter the following:

Windows

```
echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED =  
REPLACE_IMAGE_SIZE_DISABLED |  
install_path\NetBackup\bin\admincmd\bpsetconfig
```

UNIX, Linux

```
echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED =  
REPLACE_IMAGE_SIZE_DISABLED" |  
/usr/opensv/netbackup/bin/admincmd/bpsetconfig
```

Enable optimized backup of Windows deduplicated volumes

Select this attribute to enable optimized backups of optimized files on a Microsoft Data Deduplication file system.

See [“About backups and restores of Microsoft Data Deduplication file systems”](#) on page 724.

If a client has a Microsoft Data Deduplication file system configured, NetBackup performs an optimized backup of optimized files. (Unoptimized files are backed up as full, intact files.) If the client does not have a Microsoft Data Deduplication file system, a normal file backup occurs.

If you do not select this option, NetBackup still backs up the files but does so as *intact* files: the files are fully reconstructed and backed up as complete files. An unoptimized backup of optimized files is not efficient: it takes extra time and extra disk activity to reconstruct each file. However, restores are faster because NetBackup does not reconstruct the files (restores are unoptimized regardless of the type of backup).

Note: For easier management, it is recommended creating a separate policy to back up deduplicated volumes.

Only full backups are optimized. Incremental and user backups are backed up as full, intact files.

This option is supported on the Microsoft operating systems that support Data Deduplication file systems. For supported Microsoft operating systems, see the Microsoft documentation.

See [“Configuration for Microsoft Data Deduplication file system backups”](#) on page 727.

See [“Policy Attributes tab”](#) on page 663.

About backups and restores of Microsoft Data Deduplication file systems

Microsoft Data Deduplication file systems store files in one of two different states, as follows:

Unoptimized files	Files that do not qualify for <i>optimization</i> (that is, data deduplication within the file system).
Optimized files	Files that have been deconstructed, and only their unique data segments are stored on the file system.

NetBackup can do either optimized backups or unoptimized backups of the Microsoft Data Deduplication file systems, as follows:

Table 20-31 Microsoft Data Deduplication file systems backup types

Type	Description
Unoptimized backup	An unoptimized backup is one in which the Enable optimized backup of Windows deduplicated volumes policy attribute is <i>not</i> selected. NetBackup re-assembles the optimized files and backs them up as full, intact files. The storage savings of the Data Duplication file system are not retained in the backup. The optimized files remain optimized on the Data Deduplication file systems. The unoptimized files are backed up as full, intact files. NetBackup supports any type of storage destination for unoptimized backups of Microsoft Data Deduplication file systems.

Table 20-31 Microsoft Data Deduplication file systems backup types
(continued)

Type	Description
Optimized backups	An optimized backup is one in which the Enable optimized backup of Windows deduplicated volumes policy attribute is selected. Only full backups are optimized. For incremental and user backups, each file to be backed up is reconstructed and backed up in its full form. For optimized files, NetBackup backs up the chunk store and the metadata that maps the files to their segments in the chunk store. The chunk store is a location in the file system that contains the data segments that comprise the files. For unoptimized files, NetBackup backs them up as full, intact files. The following are the attributes for an optimized backup of Microsoft Data Deduplication file systems: ■ The policy must be a MS-Windows policy type. ■ The storage destination can be any Disk Storage, which is supported for Granular Backup feature. ■ The Enable Optimized Backup for supported Windows File Systems policy attribute must be selected. See “Enable optimized backup of Windows deduplicated volumes” on page 723. ■ The system drive C: cannot be backed up. Microsoft Data Deduplication file systems cannot be used for system drives. ■ Ensure that the Collect true image restore information policy attribute is not selected for this policy. TIR and optimized backups cannot be mixed. If this option is selected, the backup defaults to a non-optimized backup. The following actions are recommended for backups of Microsoft Data Deduplication file systems: ■ Run a Microsoft optimization job and a Microsoft garbage collection job before you run a full backup of deduplicated volumes. One way to do so is to configure a Deduplication Schedule in the Windows Server Manager. ■ If you want to restore entire volumes, configure a disk image backup of the volumes. A disk image backup saves time and ensures that the storage requirement can be met with the existing volume. See “Pathname rules for Windows disk image (raw) backups” on page 797. Optimized backups are supported on the Microsoft operating systems that support Data Deduplication file systems. For supported Microsoft operating systems, see the Microsoft documentation.

To restore optimized backups, use the **Restore from Optimized Backup** option in the **Backup, Archive, and Restore** interface.

Note the following items regarding restores from optimized backups:

- Can only be restored to Microsoft Data Deduplication file systems.

- Restores from optimized backups of Microsoft Data Deduplication files systems using the Microsoft deduplication API are slow. First, the image must be mounted through NFS. Then, the image must be accessed through the Microsoft deduplication API, which requests data in the amounts that may be inefficient. The process to restore an entire drive may take a very long time. Therefore, NetBackup restores all files as full, intact files (that is, non-optimized). The files are then optimized during the next scheduled file system optimization job, or you can initiate optimization manually. You should ensure that you have adequate storage for the restored files in their unoptimized state.

Configuration for Microsoft Data Deduplication file system backups

For NetBackup to do an optimized backup of Microsoft Data Deduplication file systems, specific Windows software must be installed and configured, as follows:

- Microsoft Server for NFS must be installed on the NetBackup media server. Use Windows Server Manager to add the role. Enable the NFS services. (By default, the NFS services are not enabled.)
- Microsoft Client for NFS must be installed on the Windows host that is to be backed up. Use Windows Server Manager to add the role. Without Client for NFS, restores cannot be performed.
- The Deduplication role must be installed on the Windows host that is to be backed up by using the Windows Server Manager interface.

Also, ensure that the NetBackup Client Service is configured to run as *Administrator*. If not, restores from the optimized backups fail.

Keyword phrase (policy attribute)

The **Keyword phrase** attribute is a phrase that NetBackup associates with all backups or archives based on the policy. Only the Windows and UNIX client interfaces support keyword phrases.

Clients can use the same keyword phrase for more than one policy. The same phrase for multiple policies makes it possible to link backups from related policies. For example, use the keyword phrase “legal department documents” for backups of multiple clients that require separate policies, but contain similar types of data.

The phrase can be a maximum of 128 characters in length. All printable characters are permitted including spaces and periods. By default, the keyword phrase is blank.

Clients can also specify a keyword phrase for a user backup or archive. A user keyword phrase overrides the policy phrase.

Snapshot Client and Replication Director (policy attributes)

The **Snapshot client** attributes are available when the NetBackup Enterprise Client license is installed. A snapshot is a point-in-time, read-only, disk-based copy of a client volume.

For more information about configuring snapshots, see the following guides:

- [NetBackup NAS Administrator's Guide](#)
- [NetBackup Snapshot Manager for Data Center Administrator's Guide](#)
- [NetBackup Replication Director Solutions Guide](#)
- [NetBackup for VMware Administrator's Guide](#)
- [NetBackup for Hyper-V Administrator's Guide](#)

Perform block level incremental backups (policy attributes)

The **Perform block level incremental backups** attribute allows NetBackup to back up only the changed data blocks of VMware virtual machines and Oracle or DB2 database files.

For details, refer to the appropriate NetBackup database agent guide or to the [NetBackup for VMware Administrator's Guide](#).

Use Replication Director (policy attributes)

Enable the **Use Replication Director** attribute when configuring a backup policy for Replication Director. By enabling this policy attribute, NetBackup enables other policy attributes that Replication Director requires:

- **Perform snapshot backups**
Ensures that the policy creates snapshots of the disk array.
- **Retain snapshots for Instant Recovery or SLP management**
Ensures that the policy retains the snapshot after the backup completes.
- **Perform off-host backup**
This option is selected automatically for an **NDMP** policy, along with the following selections:
 - To use **Data Mover**
 - **NDMP** as the **Machine** selection.
- **Replication Director Options**
Click the **Options** button to see the **Replication Director Options** dialog box and the default **Configuration Parameters** as follows:

Snapshot Type	■ Auto (default): The OpenStorage partner uses the best snapshot technology available to that partner to create the snapshot. ■ Differential: The OpenStorage partner creates a snapshot that is completely dependent on the source. This parameter is based on copy-on-write technology. The device creates a cache object to maintain the original blocks of the snapshot when the blocks are modified. ■ Plex: The OpenStorage Partner creates a snapshot that is completely independent of the source snapshot. This option is based on mirror-break-off technology. When a mirror device is attached to the source, the contents of the mirror device is exactly the same as the source device. When the relationship is broken between the two, the mirror device is separated from the source. The mirror device acts as a point-in-time copy. ■ Clone: The OpenStorage Partner creates an independent copy of the volume. The copy process can take some time as the entire copy must be complete. The snapshot that is created is independent of the source.
Maximum Snapshots	Sets the maximum number of Instant Recovery snapshots to be retained at one time. When the maximum is reached, snapshot rotation occurs: The next snapshot causes the oldest to be deleted. Managed by SLP retention is automatically selected if the Fixed or the Expire after Copy retention is currently selected in the SLP.

For additional information about Replication Director configuration, see the [NetBackup Replication Director Solutions Guide](#).

See “[About NetBackup Replication Director](#)” on page 1002.

Validate Policy dialog box

The **Validate Policy** dialog box appears for the backup policies that are configured for Replication Director and are not Exchange, Oracle, or VMware policies. The dialog box appears upon selecting **OK** to save and close the policy.

To ensure that the backup policy can run successfully, NetBackup validates the policy according to the validation level that you select.

Select the validation level for the policy:

- **Complete**
Performs full topology validation on underlying storage with provisioning.
Provisioning dynamically allocates NetApp storage space to each volume or LUN as data is written.
NetBackup checks the storage space on the resource pool members and performs SnapVault and SnapMirror access checks.
- **Basic**

Performs a subset of topology validation on underlying storage without provisioning.

NetBackup checks all policies to ensure that the client can perform a snapshot of the data that is indicated in the **Backup Selections** list.

NetBackup performs license checks, performs SnapVault and SnapMirror access status checks, and checks the CIFS/NFS status.

- **None**

No topology validation or provisioning.

NetBackup does not check the topology and does not provision the underlying storage.

If the policy validation finds no problems, the policy saves and closes. If validation problems are found, NetBackup displays a message that contains an error code and a description.

Perform snapshot backups (policy attributes)

The **Perform snapshot backups** attribute ensures that the policy creates snapshots of the volumes that are indicated in the policy.

Microsoft Exchange Attributes (policy attributes)

The **Microsoft Exchange** attributes let you indicate the database backup source to use for the Exchange Database Availability Group.

See the [NetBackup for Exchange Server Administrator's Guide](#).

Schedules tab

The schedules that are defined on the **Schedules** tab determine when backups occur for the selected policy. Each schedule also includes various criteria, such as how long to retain the backups.

From the policy **Schedules** tab, perform the following tasks:

- To create a new schedule, click **New**.
- To edit a schedule, select the schedule and click **Change**.
- To delete a schedule, select the schedule and click **Delete**.


Schedule attributes appear on the following tabs:

Attributes tab	Schedule the time and frequency at which a task runs, along with other scheduled attributes. See “Schedule Attributes tab” on page 731.
Start Window tab	Schedule the time of each day that a task runs. See “Start Window tab” on page 763.
Exclude Days tab	Indicate the days that a job cannot run. See “Excluding days from a schedule” on page 766.
Include Dates tab	Schedule the run days for a task by indicating specific dates, recurring weekdays, recurring days of the month. (This tab appears only when Calendar is selected as the Schedule type .) See “Include Dates tab” on page 767.

Schedule Attributes tab

The schedule **Attributes** tab contains both schedule information and other configuration options, beyond when the job is to run.

Figure 20-9 Schedule Attributes tab



The following topics describe the options on the **Attributes** tab for schedules.

Name (schedule attribute)

Specify a name for the schedule by typing it in the **Name** attribute. The schedule name appears on screens and messages about the schedule.

See [“NetBackup naming conventions”](#) on page 1054.

If the schedule is a relocation schedule created as part of a basic disk staging storage unit, the schedule name cannot be changed. The name defaults to the name of the storage unit.

See [“About staging backups”](#) on page 576.

Type of backup (schedule attribute)

The **Type of backup** attribute specifies the type of backup that the schedule controls. Select a backup type from the list. The list displays only the backup types that apply to the current policy.

If the schedule is a relocation schedule created as part of a basic disk staging storage unit, no backup type selection is needed.

[Table 20-32](#) describes the types of backups that come standard with NetBackup. [Table 20-33](#) and [Database agent backup types](#) describe other types of backups available in NetBackup.

Table 20-32 Standard backup types

Item	Description
Full Backup	Backs up all of the files that are specified in the backup selections list for the policy. The files are backed up, regardless of when the files were last modified or backed up. Full backups occur automatically according to schedule criteria. If you run incremental backups, you must also schedule a full backup to perform a complete restore. Use this option if you configure a policy for a raw partition backup (formatted partitions only).

Table 20-32 Standard backup types *(continued)*

Item	Description
Cumulative Incremental Backup	Backs up the files that are specified in the backup selections list that changed since the last full backup. All files are backed up if no previous backup was done. Cumulative incremental backups occur automatically according to schedule criteria. A complete restore requires the last full backup and the last cumulative incremental backup. Do not combine differential incremental backups and cumulative incremental backups within the same Windows policy when the incremental backups are based on archive bit (default). By default, if the time between file creation and a full or a differential incremental backup is less than 5 minutes, the differential or cumulative incremental backup may yield unexpected results. The backups are successful, but the additional files are backed up. See “About incremental backups” on page 735.
Differential Incremental Backup	Backs up the files that changed since the last successful incremental (differential or cumulative) or full backup. All files are backed up if no previous backup was done. Differential incremental backups occur automatically according to schedule criteria. A complete restore requires the last full backup, the last cumulative incremental, and all differential incremental backups that occurred since the last full backup. By default, if the time between file creation and a full or a differential incremental backup is less than 5 minutes, the differential or cumulative incremental backup may yield unexpected results. The backups are successful, but the additional files are backed up. See “About incremental backups” on page 735.
User Backup	A user initiates a user backup through the Backup, Archive, and Restore client interface. A user backup backs up all files that the user specifies. Users can start backups only during the times that are allowed on the schedule Start Window tab. For a user to be able to initiate a backup, the policy must contain a schedule of the User Backup type. Use this backup type for a catalog archive. See “Considerations for user schedules” on page 740. See “Creating a catalog archiving policy” on page 884.

Table 20-32 Standard backup types *(continued)*

Item	Description
User Archive	A user initiates a user archive through the Backup, Archive, and Restore client interface. A user archive backup first backs up the files that the user indicates. The archive then deletes the files from the local disk if the backup is successful. Archive backups free local disk space while retaining a copy for future use. The copy is kept until the retention period expires. Users can start archives only during the times that are specified in the schedule Start Window tab. For a user to be able to initiate an archive, the policy must contain a schedule of the User Archive type. Note: The NetBackup administrator should make sure that a full backup of the client exists before a user archives files from the client.

[Table 20-33](#) describes the types of backups that are available when you install NetBackup Vault.

Table 20-33 NetBackup Vault backup types

Item	Description
Automatic Vault	Applies only to Vault policies. The option does not run a backup, but instead runs the command that is specified in the Vault policy's backup selections list. In this way it starts an automatic, scheduled vault session or vault eject operation. Available only when Vault is licensed. See “Creating a Vault policy” on page 830.
Vault Catalog Backup	Use when the schedule is for a catalog backup policy that Vault uses. Available only when Vault is licensed. If this type is selected, you must configure one of the two schedule attribute combinations or the schedule cannot be saved: ■ Check and configure Multiple copies, or ■ Check Override policy storage selection, Override policy volume pool, and specify the Retention. Note: The selected storage unit selection should not be Any Available.

Database agent backup types

Each database agent often has its own set of unique backup types for a schedule. For more information, see the [NetBackup guide](#) that came with the agent.

[NetBackup documentation set](#)

About incremental backups


The following examples show how data is included in a series of full and incremental backups.

A differential incremental backup backs up the data that changed since the last full or differential incremental backup.

Note: You must run a full backup before an incremental backup. If no full backup is run, the incremental performs the role of a full backup.

Figure 20-10 shows how data is included in a series of full and differential incremental backups between January 1 and January 4.

Figure 20-10 Full and differential incremental example



The January 1 full backup includes all files and directories in the policy backup selections list. The subsequent differential incremental backups include only the data that changed since the last full or differential incremental backup. If the disk fails sometime on January 4 (after the backup), the full backup and all three of the incremental backups are required for the recovery.

A cumulative incremental backup backs up the data that changed since the last full backup. Figure 20-11 shows how data is included in a series of full and cumulative incremental backups between January 1 and January 4. The January 1 full backup includes all files and directories in the policy backup selections list. Each of the cumulative incremental backups includes the data that changed since the last full backup. If the disk fails sometime on January 4 (after the backup), the full backup and the last cumulative incremental backup are required for the recovery.

Figure 20-11 Full and cumulative incremental example




Table 20-34 describes how to determine the retention of differential and cumulative incremental backups to prevent a gap in backup coverage.

Table 20-34 Retention requirements for incremental backups

Type	Retention requirement	Comments
Differential	Longer	To restore all files requires the last full backup and all the differential incremental backups that occurred since the last full backup. Therefore, all the differentials must be kept until the next full backup occurs.
Cumulative	Shorter	Each cumulative incremental backup contains all the changes that occurred since the last full backup. Therefore, a complete restore requires only the most recent cumulative incremental in addition to the full backup.

Table 20-35 compares the advantages and disadvantages of using differential or cumulative incremental backups based on possible backup and restore times.

Table 20-35 Relative backup and restore times for incremental backups

Type	Backup time	Restore time	Comments
Differential	Shorter	Longer	Less data in each backup, but all differential incremental backups are required since the last full backup for a restore. This results in a longer restore time.
Cumulative	Longer	Shorter	More data in each backup, but only the last cumulative incremental backup is required for a complete restore (in addition to the full).

You can use a combination of cumulative and differential incremental backups together to get the advantages of both methods. For example, assume a set of schedules with the following backup frequencies and retention periods. (Notice that the differential incremental backups occur more often.)

Table 20-36
 Example frequencies and retention periods

Backup type	Frequency	Retention period
Full	6 days	2 weeks
Cumulative incremental	2 days	4 days
Differential incremental	1 day	2 days

The schedules that are described in [Table 20-36](#) result in the following series of backups:

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8
Full	Diff	Cum	Diff	Cum	Diff	Full	Diff

The example produces the following results:

- Every other day a differential incremental backup occurs, which usually has a minimum backup time.
- On alternate days, a cumulative incremental backup occurs, which requires more time than the differential backup, but not as much time as a full backup. The differential backup can now be expired.
- To recover all files may require (at most), two incremental backups in addition to the most recent full backup. The combination of backups usually means less restore time than if all differential incremental backups were used. The full backups can be done less often if the amount of data being backed up by the incremental backups is small.

How NetBackup determines when Windows files are due for backup

On Windows clients, NetBackup performs the incremental backups when the **Perform incrementals based on archive bit** setting is enabled. This setting is found in the **Backup, Archive, and Restore** client interface, under **File > NetBackup Client Properties**, on the **General** tab.

If **Perform incrementals based on archive bit** is enabled, incremental backups for the client are based on the state of the archive bit of each file. The operating system sets the bit whenever a file changes, and it remains set until cleared by NetBackup. The conditions under which NetBackup clears the bit depend on the type of backup being performed.

Full Backup	NetBackup backs up files regardless of the state of their archive bit. After a full backup, the archive bit is always cleared.
Differential Incremental Backup	NetBackup backs up the files that have the archive bit set and have therefore changed. When the client receives a response from the server that indicates that the backup was successful (or partially successful) the archive bits are cleared. The clear archive bit lets the next differential incremental backup back up only the files that changed since the previous full or differential incremental backup.
Cumulative Incremental Backup	NetBackup backs up the files that have the archive bit set. However, NetBackup does not clear the archive bits after the backup. Without a clear archive bit, the next cumulative incremental backup backs up changed files and the files that were in the cumulative incremental backup.

If **Perform incrementals based on archive bit** is disabled, NetBackup includes a file in an incremental backup only if the datetime stamp of the file has changed since the last backup. The datetime stamp indicates when the file was last backed up. The backup types use the datetime stamp differently.

Full Backup	NetBackup backs up files regardless of the datetime stamp.
Differential Incremental Backup	NetBackup compares the datetime stamp of the file against the last full or incremental backup.
Cumulative Incremental Backup	NetBackup compares the datetime stamp of the file against the last full backup.

If files are installed or copied from another computer, the new files retain the datetime stamp of the originals. If the original date is before the last backup date, the new files are not backed up until the next full backup.

How NetBackup determines when UNIX files are due for backup

Incremental backups on UNIX clients consider all files and directories to determine if a backup is due based on a reference date. (That is, back up all the files that have changed since *date_x*).

The following types of time are associated with UNIX files and directories:

<code>mtime</code>	The file modification time. The file system updates the <code>mtime</code> for a file or directory each time the file is modified. An application can save the <code>mtime</code> of the file before it modifies it. The application then resets it with the <code>utime(2)</code> system call.
<code>atime</code>	The file access time. The file system updates the <code>atime</code> for a file or directory each time the file is accessed (read or write). An application can save the <code>atime</code> of the file before it accesses it. The application then resets it with the <code>utime(2)</code> system call.
<code>ctime</code>	The inode change time. The <code>ctime</code> for a file or directory is updated each time the file or directory's inode changes. (For example, changes due to permissions, ownership, and link-counts changes.) The <code>ctime</code> for a file or directory cannot be saved before a change, and then reset after a change. The <code>ctime</code> of a file or a directory changes when the <code>mtime</code> and <code>atime</code> (changes with the <code>utime(2)</code> system call) is reset.

When NetBackup reads the data for a file that is included in a backup, it does not affect the file modification time. It does affect the access time of the file. For this reason, NetBackup saves the `atime` and `mtime` of the file before it reads the file. Then NetBackup resets the `atime` and `mtime` with the `utime(2)` system call. NetBackup does not cause problems for storage migration products or the administrator scripts that use file access times (`atime`) as criteria for their operations. While this benefit is obvious, a side effect is that it does update the `ctime` of the file.

Note: When NetBackup Accelerator is used to perform backups, NetBackup does not reset the `atime` after the file is backed up. This avoids updating the `ctime`, because a `ctime` update would require a loss in Accelerator optimization.

See [“Accelerator notes and requirements”](#) on page 707.

Customers can configure NetBackup so that it does not reset the access time of the file after it reads a file. Customers can choose to have NetBackup use the `ctime` and the `mtime` of the file to determine what files to include in an incremental backup. Normally, these two options are used together, but there may be some sites that want to use one without the other. By default, NetBackup uses only the `mtime` of the file to determine what files and directories to back up.

When a file is moved from one location to another, the `ctime` of the file changes, but the `mtime` remains unchanged. If NetBackup uses only the `mtime` to determine the files that are due during an incremental backup, it does not detect these moved files. For sites where using the `mtime` might create a problem, use the `ctime` to determine files due to be included in an incremental backup. The `ctime` is used if

the `bp.conf` file contains the `USE_CTIME_FOR_INCREMENTALS` and `DO_NOT_RESET_FILE_ACCESS_TIME` entries.

When a directory is moved from one location to another, the `ctime` of the directory changes, but the `mtime` remains unchanged. Neither the `mtime` nor the `ctime` are changed for the files or directories within the moved directory. No reliable method using file timestamps can determine that files within a moved directory need to be included in an incremental backup.

In either case, these moved files and directories are included in subsequent full backups.

Considerations for user schedules

In order for users to perform backups and archives, an administrator must create a schedule that allows user backups.

User backup schedules and user archive schedules can be included in a policy that contains automatic backup schedules. If you create separate policies for user backups or user archives, the considerations are similar to those for automatic backups. In user backup schedules, however, no backup selection list is necessary because users select the objects before they start the backup or archive.

To use a specific policy or schedule for user backups or user archives, perform the tasks that are specified for each client type:

Table 20-37 Tasks for specifying a policy or schedule for user backups or user archives

Client type	Task
Microsoft Windows clients	- Start the Backup, Archive, and Restore client interface. - On the File menu, click NetBackup Client Properties - Select the Backups tab, and specify the backup policy and backup schedule.
UNIX clients	Specify the policy and schedule with <code>BPARCHIVE_POLICY</code> , <code>BPARCHIVE_SCHED</code> , <code>BPBACKUP_POLICY</code> , or <code>BPBACKUP_SCHED</code> options in the <code>bp.conf</code> file.

Restores can be performed at any time and are not scheduled.

Note: An archive is different from a backup. During an archive, NetBackup first backs up the selected files, and then deletes the files from the local disk if the backup is successful. In this topic, references to backups also apply to the backup portion of archive operations unless otherwise noted.

How to plan schedules for user backups and user archives

To plan schedules for user backups and user archives, consider the following:

Automatic backups	If possible, do not permit user backups and user archives when automatic backups are running. If an automatic backup is running when a user submits a backup or archive, NetBackup usually queues the user job. The job is not queued if there is a limiting setting. (For example, the Limit jobs per policy policy attribute or the Maximum jobs per client Global attributes host property.) See “Limit jobs per policy (policy attribute)” on page 678. See “Global attributes properties” on page 109. If the automatic backup continues to run, the user job may miss the backup window depending on how the limiting settings are configured. On the other hand, user jobs can delay automatic backups and can cause backups to miss the backup window.
Storage units	Use a different storage unit to eliminate conflicts with automatic backups.
Volume pools	Use a different volume pool to manage the media separate from the automatic backup media.
Retention periods	Consider setting the retention period for archives to infinite, since the disk copy of the files is deleted. Note: If the retention period expires for a backup, it can be difficult or impossible to restore the archives or backups.

Synthetic backup (schedule attribute)

The **Synthetic backup** schedule attribute allows a backup to be assembled from previous backups. A synthetic backup may be a synthetic full or a synthetic cumulative incremental backup. The backups include one previous, traditional full backup, and subsequent differential backups, and a cumulative incremental backup. (A traditional full backup means a non-synthesized, full backup.) A client can then use the synthesized backup to restore files and directories in the same way that a client restores from a traditional backup.

Synthetic backups can be written to tape, to disk storage units, or to a combination of both.

See [“About synthetic backups”](#) on page 842.

Accelerator forced rescan option (schedule attribute)

The policy **Schedules** tab contains an option called **Accelerator forced rescan**. This option creates a checksum of the content of each file during backup. It uses the checksums for change detection. It provides a safety net by establishing a new baseline for the next Accelerator backup.

Note: The following information is only applicable to Files and Folder (Unix/Windows) backups and NDMP backups. The Accelerator forced rescan schedule is not necessary for other backup types such as VMware and Hyper-V.

The **Accelerator forced rescan** option detects the following events:

- The file's data changes but the file's metadata does not change.
- The file's metadata becomes corrupted such that it does not indicate that the file has changed.
- A malicious user or application changes the file's metadata such that it does not indicate that the file has changed.

Note: If **Accelerator forced rescan** is enabled, NetBackup uses the Windows change journal to help determine if a file has changed. NetBackup also examines the file system metadata for each file to help detect which files have changed.

For the most efficient use of Accelerator, Accelerator policies must include at least two full-backup schedules: one full schedule with the **Accelerator forced rescan** option disabled, and another full schedule with **Accelerator forced rescan** enabled. See [Table 20-38](#).

Table 20-38 Required full-backup schedules for each Accelerator policy

Full backup schedules	Notes on schedule frequency
First schedule: Accelerator forced rescan disabled	Configure this schedule to run most of your Accelerator full backups.

Table 20-38 Required full-backup schedules for each Accelerator policy
(continued)

Full backup schedules	Notes on schedule frequency
Second schedule: Accelerator forced rescan enabled	Configure this schedule to run less often than the first full-backup schedule. For example: If the first full-backup schedule runs weekly, run the second schedule (with the Accelerator forced rescan option enabled) every few months. However, the best frequency for this schedule depends upon your environment. Note: If the policy has no schedule that enables the Accelerator forced rescan option, all full backups automatically enable that option and backup performance is reduced.

Note the following about the **Accelerator forced rescan** option:

- The **Accelerator forced rescan** option is grayed out if the **Use Accelerator** option on the **Attributes** tab is not selected.
- Because of the checksum processing on the client, this option reduces backup speed as compared to the **Use Accelerator** option on its own. The speed reduction depends on the client's configuration and its current processing load. If the client is busy with many jobs when Accelerator backup begins, checksum processing can reduce backup speed.
- If the Windows volume is not NTFS or ReFS, and the volume has no full backup schedule that is configured with the **Accelerator forced rescan** option, note: NetBackup uses **Accelerator forced rescan** on every full backup. The **Accelerator forced rescan** option is not enabled in the schedule but is in fact used, as indicated in the NetBackup log messages. This option may reduce the speed of the backup as compared to the previous backups that did not use **Accelerator forced rescan**.
See [Table 20-38](#) for recommended full-backup schedules.
- For an MS-Windows or Standard policy (to back up files and folders), all the data on the client is read. However, NetBackup sends only the changed data over the network to be included in the backup image. Sending only the changed data is similar to a regular Accelerator full backup. Thus, for an Accelerator forced rescan backup, the optimization percentage is similar to a regular Accelerator full backup. Note that the duration of the backup falls somewhere between a non-Accelerator full backup and a regular Accelerator full backup.

Calendar (schedule attribute)

Calendar-based schedules allow administrators to create a job schedule based on a calendar view. Select **Calendar** to display the **Include Dates** tab.

See [“Include Dates tab”](#) on page 767.

A calendar-based relocation schedule determines the days that images are swept from the disk staging storage unit to the final destination storage unit. (A relocation schedule is created as part of a basic disk staging storage unit configuration.)

Enable **Retries allowed after runday** to have NetBackup attempt to complete the schedule until the backup is successful. With this attribute enabled, the schedule attempts to run, even after a specified run day has passed.

Frequency (schedule attribute)

Use the **Frequency** attribute to specify how much time must elapse between the successful completion of a scheduled task and the next attempt.

For example, assume that a schedule is set up for a full backup with a frequency of one week. If NetBackup successfully completes a full backup for all clients on Monday, it does not attempt another backup for this schedule until the following Monday.

To set the frequency, select a frequency value from the list. The frequency can be seconds, minutes, hours, days, or weeks.

A frequency-based relocation schedule determines how often images are swept from the basic disk staging storage unit to the final destination storage unit. (A relocation schedule is created as part of a basic disk staging storage unit configuration.)

NetBackup recognizes the intervals that suggest schedules based on days, even if the job does not run daily. For example, if the frequency is 48 hours, NetBackup tries to run the job at the same time every other day. (NetBackup checks if the frequency is divisible by 24 hours.) If the interval is not divisible by 24, NetBackup does not attempt to run the job at about the same time of day. Instead, NetBackup tries to run the job at the indicated interval after the last successful backup. (For example, 52 hours later.)

Note: **Frequency** does not apply to user schedules because the user can perform a backup or archive whenever the time window is open.

About backup frequency

To determine backup frequency, consider how often data changes. For example, determine if files change several times a day, once a day, weekly, or monthly.

Typically, sites perform daily backups to preserve daily work. Daily backups ensure that only one day's work is lost in case of a disk failure. More frequent backups are necessary when important data changes many times during the day and the changes would be difficult to reconstruct.

Daily backups are usually the incremental backups that record the changes since the last incremental or full backup. Incremental backups conserve resources because they use less storage and take less time to perform than full backups.

Full backups usually occur less frequently than incremental backups but should occur often enough to avoid accumulating consecutive incremental backups. A large number of incremental backups between full backups increases the time it takes to restore a file. The time increases because of the effort that is required to merge the incremental backups when files and directories upon restore.

Consider the following when setting the frequency for full backups:

- Extend the time between full backups for the files that seldom change. A longer frequency uses fewer system resources. It also does not significantly increase recovery time because the incremental backups between full backups are smaller.
- Decrease the time between full backups for the files that change frequently. A shorter frequency decreases restore time. A shorter time between full backups can also use fewer resources. It reduces the cumulative effect of the longer incremental backups that are necessary to keep up with frequent changes in the files.

To achieve the most efficient use of resources, ensure that most of the files in a given policy change at about the same rate. For example, assume that half of the files in a policy selection list change frequently enough to require a full backup every week. However, the remaining files seldom change and require monthly full backups only. If all the files are in the same policy, full backups are performed weekly on all the files. This wastes system resources because half the files need full backups only once a month. A better approach is to divide the backups into two policies, each with the appropriate backup schedule, or to use synthetic backups.

If more than one automatic schedule is due for a client within a policy, the backup frequency determines the schedule that NetBackup uses as follows:

- Jobs from the schedule with the lower frequency (longer period between backups) always have higher priority. For example, a schedule that has a backup frequency of one month takes priority over a schedule with a backup frequency of 2 weeks.

- When two schedules are each due to run, the schedule with the schedule name that is first in alphabetical order runs first. Alphabetical priority occurs if both of the following are true:
 - Each schedule is within the defined time window.
 - Each schedule is configured with the same frequency value.

NetBackup prioritizes the example schedules in the following order:

Table 20-39 Examples of schedule frequency and priority

Schedule Name	Frequency	Priority
monthly_full	One month	First
weekly_full	One week	Second
daily_incremental	One day	Third

Instant Recovery (schedule attribute)

The **Instant Recovery** attributes are available under the following conditions:

- The **Snapshot Client** option is licensed and installed.
Refer to the [NetBackup NAS Administrator's Guide](#) and the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#).
- **Perform snapshot backups** is selected.
- **Retain snapshots for Instant Recovery** is selected.

See “[Snapshot Client and Replication Director \(policy attributes\)](#)” on page 728.

This attribute has two options.

Snapshots and copy snapshots to a storage unit

The snapshot persists on the client volume with a backup copy made to the storage unit on the media server.

Snapshots only

The snapshot is not backed up to tape or to other storage. NetBackup creates a snapshot on disk only. This option is required for the **NAS_Snapshot** method.

The snapshot is created on the same device as the one that contains the original data if it uses **VxFS_Checkpoint** method or is VxVM space optimized. In this case, another policy can be used to back up the data to a separate device.

Transaction logs are not truncated at the end of the backup.

The **Instant Recovery** attributes are grayed out if the **Policy storage** option on the policy **Attributes** tab refers to a storage lifecycle policy. If that is the case, the storage lifecycle policy configuration governs the **Instant Recovery** attributes.

However, the **Override policy storage selection** attribute on the Schedule **Attributes** tab overrides the **Policy storage** option. If a storage unit is selected on the Schedule **Attributes** tab, the **Instant Recovery** attributes become enabled.

See [“Policy storage \(policy attribute\)”](#) on page 669.

See [“Override policy storage \(schedule attribute\)”](#) on page 751.

Multiple copies (schedule attribute)

When the **Multiple copies** attribute is enabled, NetBackup can create up to four copies of a backup simultaneously. The storage units must be on the same media server with sufficient resources available for each copy. For example, to create four copies simultaneously in a Media Manager storage unit, the unit needs four tape drives. (This option is sometimes referred to as Inline Copy, Inline Tape Copy, or ITC.)

To create more than four copies, additional copies can be created at a later time using duplication.

If multiple original images are created simultaneously, the backup time that is required may be longer than for one copy. Also, if both Media Manager and disk storage units are specified, the duration of disk write operations match that of slower removable media write operations.

About configuring multiple copies

To create multiple copies, the following criteria must be met:

- The backup destinations must share the same media server with sufficient resources available for each copy.
- The storage units that are used for multiple copies must be configured to allow a sufficient number of concurrent jobs to support the concurrent copies. The pertinent storage unit settings are **Maximum concurrent jobs** and **Maximum concurrent write drives**.

See [“Maximum concurrent jobs storage unit setting”](#) on page 563.

See [“Maximum concurrent write drives storage unit setting”](#) on page 563.

Multiple copy operations do not support the following:

- NDMP storage units
- Synthetic backups
- Storage lifecycle policies

Storage lifecycle policies offer their own method to create multiple copies.
See [“About writing multiple copies using a storage lifecycle policy”](#) on page 639.

Configure Multiple Copies dialog box

The **Configure Multiple Copies** dialog box contains the following options:

Table 20-40 Configure Multiple Copies dialog box

Field	Description
Copies	NetBackup can create up to four copies of a backup simultaneously. The storage units must be on the same media server and there must be sufficient resources available for each copy. To create more than 4 copies, create additional copies at a later time by using duplication.
Primary copy	Copy 1 is the primary copy. If Copy 1 fails for some reason, the first successful copy is the primary copy. See “Promoting a copy to a primary copy” on page 927.
Storage unit	Specify the storage unit where each copy is to be stored. If a Media Manager storage unit has multiple drives, you can use it for both the source and the destination. To let NetBackup decide at run-time, select Any Available .
Volume pool	Indicate where each copy is to be stored.
Retention schedule	Specify how long NetBackup retains the backups. See “Retention (schedule attribute)” on page 753.
If this copy fails	In the event that the copy does not complete, select whether you want the entire job to fail (fail all copies), or whether you want the remaining copies to continue. Regardless of how the fail or continue flag is set, all the copy jobs wait in the queue until resources are available for all copies. The first job does not start until the copies have resources. If a copy is configured to allow other copies to continue the job if the copy fails, and if Checkpoint restart for backup jobs is selected for this policy, only the last failed copy that contains a checkpoint can be resumed. See “Take checkpoints every __ minutes (policy attribute)” on page 674.

Table 20-40 Configure Multiple Copies dialog box (*continued*)

Field	Description
Media owner	Select who should own the media onto which NetBackup writes the images. The following options are available: ■ Any Lets NetBackup select the media owner, either a media server or server group. ■ None Specifies that the media server that writes to the media that owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ A server group Specify a media server group to allow only those media servers in the group to write to the media on which backup images for this policy are written. All media server groups that are configured in the NetBackup environment appear in the drop-down list. See “Add a server group” on page 374.

Configuring multiple copies in a policy schedule

To configure a policy schedule to create multiple copies, use the following procedure.

To configure a schedule to create multiple copies

1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.

2 Do one of the following:

- To change an existing policy
- Select the policy to change.
 - On the **Edit** menu, click **Change**.

- To create a new policy
- On the **Actions** menu, click **New > Policy**.
 - Name the policy, and click **OK**.

3 Select the **Schedules** tab.

4 Double-click an existing schedule or click **New** to create a new schedule.

5 In the dialog box that appears, click in the **Attributes** tab, select **Multiple copies**, and then click **Configure**.

If the destination for this policy is a storage lifecycle policy, the **Multiple copies** box is unchecked. NetBackup does not allow the two methods for creating multiple copies to be enabled at the same time.

See [“Policy storage \(policy attribute\)”](#) on page 669.

See [“About writing multiple copies using a storage lifecycle policy”](#) on page 639.

- 6
- In the **Copies** field, specify the number of copies to be created simultaneously. The number must be between 1 and 4.

Configure Multiple Copies

Copies: 2 All storage units must be connected to the same media server.

	Primary:	Storage unit:	Volume pool:	Retention:	For each image if this copy fails:	Media Owner:
Copy 1:	✓	ADV_DISK-master-stu	NetBackup	2 weeks (Retention ...	continue	Any
Copy 2:		MSDP-media-stu	NetBackup	2 months (Retention...	continue	Any
Copy 3:						
Copy 4:						

OKCancelHelp

Copy 1 is the primary copy. If **Copy 1** fails, the first successful copy is the primary copy.

Usually, NetBackup restores from the primary copy of an image. However, it is possible to restore from a specific backup copy other than the primary copy. To do so, use the `bprestore` command.

To create more than four copies, create additional copies at a later time by using duplication.

See “[Configure Multiple Copies dialog box](#)” on page 748.

See “[About configuring multiple copies](#)” on page 747.

- 7
- Specify the storage unit where each copy is stored. Select **Any Available** to allow NetBackup to select the storage unit at run-time.

If a Media Manager storage unit contains multiple drives, the storage unit can be used for both the original image and the copies.

- 8
- Specify the volume pool where each copy is stored.

- 9
- Select the retention level for each copy.

See “[Retention \(schedule attribute\)](#)” on page 753.

10 Select one of the following from the **If this copy fails** list:

- continue

Continues making the remaining copies.
Note: If **Take checkpoints every __ minutes** is selected for this policy, only the last failed copy that contains a checkpoint can be resumed.

See [“Take checkpoints every __ minutes \(policy attribute\)”](#) on page 674.
- fail all copies

Fails the entire job.

11 For tape media, specify who should own the media onto which NetBackup writes the images:

- Any

NetBackup selects the media owner, either a media server or server group.
- None

Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
- A server group

Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All the media server groups that are configured in the NetBackup environment appear in the list.

These settings do not affect images residing on disk. One media server does not own the images that reside on shared disks. Any media server with access to the shared pool of disk can access the images.

12 Click **OK** until the policy is saved.

Override policy storage (schedule attribute)

The **Override policy storage selection** attribute works as follows:

- Disabled

Instructs the schedule to use the **Policy storage** as specified on the policy **Attributes** tab.
- Enabled

Instructs the schedule to override the **Policy storage** as specified on the policy **Attributes** tab.

Select the storage from the list of previously configured storage units and storage lifecycle policies. If the list is empty, no storage has been configured.

See [“Policy storage \(policy attribute\)”](#) on page 669.

If a data classification is indicated for the policy, only those storage lifecycles with the same data classification appear in the list.

See [“Data classifications \(policy attribute\)”](#) on page 668.

Note: Storage lifecycle policies cannot be selected within the **Configure Multiple Copies** dialog box.

See [“About configuring multiple copies”](#) on page 747.

Override policy volume pool (schedule attribute)

The **Override policy volume pool** attribute works as follows:

Disabled	Instructs the schedule to override the volume pool that is specified as the Policy volume pool on the policy Attribute tab. If no policy volume pool is specified, NetBackup uses NetBackup as the default. If the policy is for a NetBackup catalog, NBU-Catalog policies use CatalogBackup .
Enabled	Instructs the schedule to override the volume pool that is specified as the Policy volume pool on the policy Attribute tab. Select the volume pool from the list of previously configured volume pools.

See [“Policy volume pool \(policy attribute\)”](#) on page 672.

Override media owner (schedule attribute)

The **Override media owner** attribute applies only to tape media. It specifies whether to use the policy media owner or another owner for the schedule. The rules for shared disk media are more flexible so override settings are not needed for disk media.

The attribute works as follows:

Disabled	Instructs the schedule to use the media owner that is specified as the Media Owner in the policy Attribute tab.

Enabled	Instructs the schedule to override the media owner that is specified as the Media Owner in the policy Attribute tab. Select the new media owner from the list: ■ Any. NetBackup selects the media owner, either a media server or server group ■ None Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ A server group Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All media server groups that are configured in the NetBackup environment appear in the list.

See [“Media Owner \(policy attribute\)”](#) on page 681.

Retention (schedule attribute)

The **Retention** attribute specifies how long NetBackup retains the backups. To set the retention period, select a time period (or level) from the list. When the retention period expires, NetBackup deletes information about the expired backup. After the backup expires, the files in the backup are unavailable for restores. For example, if the retention is 2 weeks, data can be restored from a backup that this schedule performs for only 2 weeks after the backup.

If a policy is configured to back up to a storage lifecycle policy, the **Retention** attribute in the schedule is ignored. The retention period that the lifecycle indicates is followed instead.

See [“Creating a storage lifecycle policy”](#) on page 602.

See [“Retention periods properties”](#) on page 152.

About assigning retention periods

The retention period for data depends on the likelihood of restoring information from media after a certain period of time. Some types of data (financial records, for example) have legal requirements that determine the retention level. Other data (preliminary documents, for example) can probably be expired when the final version is complete.

A backup’s retention also depends on what needs to be recovered from the backup. For example, if day-to-day changes are critical, keep all the incremental backups

in addition to the full backups for as long as the data is needed. If incremental backups only track work in progress toward monthly reports, expire the incremental backups sooner. Rely on the full backups for long-term recovery.

Establish some guidelines that apply to most of the data to determine retention periods. Note the files or the directories that have retention requirements outside of these guidelines. Plan to create separate policies for the data that falls outside of the retention requirement guidelines. For example, place the files and directories with longer retention requirements in a separate policy. Schedule longer retention times for the separate policies without keeping all policies for the longer retention period.

The following table describes recommended retention periods for different types of backups.

Table 20-41 Recommended retention periods for different types of backups

Type of backup	Description
Full Backup	Specify a time period that is longer than the frequency setting for the schedule. (The frequency is how often the backup runs). For example, if the frequency is one week, specify a retention period of 2-4 weeks. Two to 4 weeks provides enough of a margin to ensure that the current full backup does not expire before the next full backup occurs.
Differential Incremental Backup	Specify a time period that is longer than the period between full backups. For example, if full backups occur weekly, save the incremental backups for 2 weeks.
Cumulative Incremental Backup	Specify a time period that is longer than the frequency setting for the schedule. (The frequency is how often the backup runs). For example, if the frequency setting is one day, specify a retention period of one week. One week provides enough of a margin to ensure that the current cumulative-incremental backup does not expire before the next successful one occurs. A complete restore requires the previous full backup plus the most recent cumulative-incremental backup.

The following table suggests several ways that you can prevent backups from expiring earlier than desired.

Table 20-42 Suggestions for preventing prematurely expired backups

Item	Description
Retention period	Assign an adequate retention period. NetBackup does not track backups after the retention period expires. Recovering files is difficult or impossible after the retention period expires. For the backups that must be kept for more than one year, set the retention period to infinite.
Full backups and incremental backups	Assign a longer retention period to full backups than to incremental backups within a policy. A complete restore requires the previous full backup plus all subsequent incremental backups. It may not be possible to restore all the files if the full backup expires before the incremental backups.
Archive schedules	Set the retention period to infinite.
Tape	Set the retention period to infinite. If infinite is unacceptable because of NetBackup database space limitations, set the retention period to match the length of time that the data is to be retained.

Another consideration for data retention is off-site storage of the backup media. Off-site storage protects against the disasters that may occur at the primary site. Consider the following off-site storage methods as precautions for disaster recovery:

- Use the duplication feature to make a second copy for off-site storage.
- Send monthly or weekly full backups to an off-site storage facility.
To restore the data, request the media from the facility. To restore a total directory or disk with incremental backups requires the last full backup plus all incremental backups.
- Configure an extra set of schedules to create the backups to use as duplicates for off-site storage.

Regardless of the method that is used for off-site storage, ensure that adequate retention periods are configured. Use the NetBackup import feature to retrieve expired backups.

By default, NetBackup stores each backup on a tape volume that contains existing backups at the same retention level. If a backup has a retention level of 2, NetBackup stores it on a tape volume with other backups at retention level 2. When NetBackup encounters a backup with a different retention level, it switches to an appropriate volume. Because tape volumes remain assigned to NetBackup until all the backups on the tape expire, this approach results in more efficient use of media. One small

backup with an infinite retention prevents a volume from being reused, even if all other backups on the volume expired.

To mix retention levels on volumes, select **Allow multiple retentions per media** in the **Media** host properties.

If you keep only one retention level on each volume, do not use any more retention levels than necessary. Multiple retention levels increase the number of required volumes.

See [“Media properties”](#) on page 119.

Note: Retention levels can be mixed on disk volumes with no restrictions.

See [“Changing a retention period”](#) on page 154.

Media multiplexing (schedule attribute)

The **Media multiplexing** attribute specifies the maximum number of jobs from the schedule that NetBackup can multiplex onto any one drive. Multiplexing sends concurrent backup jobs from one or several clients to a single drive and multiplexes the backups onto the media.

Specify a number from 1 through 32, where 1 specifies no multiplexing. Any changes take effect the next time a schedule runs.

Note: Some policy types and some schedule types do not support media multiplexing. The option cannot be selected in those instances.


Caution: For MS-SQL-Server policies, do not enable multiplexing for a schedule that is also configured to backup with multiple stripes. Restores fail when multiplexing is enabled for a schedule that uses more than one stripe.

To configure multiplexed backups, multiplexing must be indicated in both the storage unit (**Maximum streams per drive** setting) and the schedule (**Media multiplexing** setting) configuration. Regardless of the **Media multiplexing** setting, the maximum number of jobs that NetBackup starts never exceeds the **Maximum streams per drive** value for the storage unit.

NetBackup multiplexing sends concurrent backups from one or several clients to a single storage device. NetBackup multiplexes the backups sequentially onto the media. Multiplexed and unmultiplexed backups can reside on the same volume. Separate volume pools or media IDs are not necessary.

Figure 20-12 shows the multiplexed flow of client data to a server.

Figure 20-12 Multiplexed backups



About multiplexing

Multiplexing is generally used to reduce the amount of time that is required to complete backups. The following table describes circumstances where performance improves by using multiplexing:

Table 20-43 Circumstances where multiplexing improves performance

Item	Description
Slow clients	Instances in which NetBackup uses software compression, which normally reduces client performance, are also improved.
Multiple slow networks	The parallel data streams take advantage of whatever network capacity is available.
Many short backups (for example, incremental backups)	In addition to providing parallel data streams, multiplexing reduces the time each job waits for a device to become available. Therefore, the storage device transfer rate is maximized.

No special action is required to restore a multiplexed backup. NetBackup finds the media and restores the requested backup. Multiplexing reduces performance on restores because it uses extra time to read the images.

To reduce the effect of multiplexing on restore times, set the storage unit maximum fragment size to a value smaller than the largest allowed value. Also, on UNIX, enable fast-tape positioning (locate block), if it applies to the tape drives in use.

Consider the following configuration settings when using multiplexing.

Table 20-44 Properties and attributes that affect multiplexing

Item	Description	Where to find it
Limit jobs per policy (policy attribute)	Limits the number of jobs that NetBackup performs concurrently when a policy is run. Set this attribute high enough to support the specified level of multiplexing. See "Limit jobs per policy (policy attribute)" on page 678.	■ In the policy attributes.

Table 20-44 Properties and attributes that affect multiplexing (*continued*)

Item	Description	Where to find it
Maximum jobs per client (host property)	Limits the number of backup jobs that can run concurrently on any NetBackup client. This property is part of Global attributes host properties. See “Global attributes properties” on page 109. Usually, the client setting does not affect multiplexing. However, consider a case where jobs from different schedules on the same client go to the same storage unit. In this case, the maximum number of jobs that are permitted on the client is reached before the multiplexing limit is reached for the storage unit. When the maximum number of jobs on the client is reached, NetBackup cannot use the storage unit’s full multiplexing capabilities. Select a value that is based on the ability of the central processing unit to handle parallel jobs. Because extra buffers are required, memory is also important. If the server cannot perform other tasks or runs out of memory or processes, reduce the Maximum streams per drive setting for the storage unit. To estimate the potential load that multiplexing can place on the central processing unit, consider the following limits: ■ The number of maximum concurrent jobs that NetBackup can attempt equals the sum of the concurrent backup jobs that can run on all storage units. ■ The maximum concurrent jobs that can run on a storage unit equals the value of Maximum streams per drive, multiplied by the number of drives. See “Maximum streams per drive storage unit setting” on page 565.	■ Hosts > Host properties > Select the primary server > Global Attributes.


Table 20-44 Properties and attributes that affect multiplexing (*continued*)

Item	Description	Where to find it
Maximum data streams (host property)	Set the maximum number of jobs that are allowed on a specific client without affecting other clients. This property is part of Client Attributes host properties. See “General tab of the Client attributes properties” on page 66.	■ Hosts > Host properties > Select the primary server > Client attributes. ■ The property appears in the right pane on the General tab.
Delay on multiplexed restores (host property)	Specifies how long the server waits for additional restore requests of files and raw partitions in a set of multiplexed images on the same tape. This property is part of the General server host properties. See “General tab of the Client attributes properties” on page 66.	■ Hosts > Host properties > Select the primary server > General server.
Media Multiplexing (policy schedule attribute)	If the limit is reached for a drive, NetBackup sends jobs to other drives. When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches the Media Multiplexing limit or the Maximum streams per drive limit. See “Media multiplexing (schedule attribute)” on page 756.	■ Open the policy. ■ Select the Schedules tab. ■ Select Add to create a new schedule and configure the Media multiplexing option.
Maximum streams per drive (storage unit setting)	NetBackup can add jobs from more than one schedule to a drive. When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches the Maximum streams per drive limit or the Media multiplexing limit. See “Maximum streams per drive storage unit setting” on page 565.	■ Select Storage > Storage units. ■ Open the storage unit. ■ The setting is available in the Basic properties for the storage unit.

Example of using multiplexing with schedules

[Figure 20-13](#) provides an example of how schedules are affected when multiplexing is active.

Figure 20-13 Multiplexing process scenario



Assume the following about Figure 20-13.

- Schedule A begins first.
Schedules can be in the same or in different policies.
- **Allow Multiple Data Streams** is enabled.
Consequently, a client can have multiple data streams.
See “[Allow multiple data streams \(policy attribute\)](#)” on page 697.

Table 20-45 Description of the multiplexing process scenario

Event	Description
1 and 2	■ Jobs A1 and A2 from client <i>DOG</i> start on Drive 1. ■ For Schedule A, the Media Multiplexing limit of 2 is reached for Drive 1.
3 and 4	■ Jobs A3 and A4 from client <i>CAT</i> start on Drive 2. ■ For Schedule A, the Media Multiplexing limit of 2 is reached for Drive 2.
5 and 6	■ Jobs B1 and B2 for client <i>FOX</i> start on Drive 1. ■ The Maximum streams per drive storage unit setting is reached for Drive 1.

Table 20-45
Description of the multiplexing process scenario *(continued)*

Event	Description
7 and 8	- Jobs B3 and B4 from client <i>OTTER</i> start on Drive 2. - All jobs are now running for Schedule B. - The Maximum streams per drive storage unit setting is reached for Drive 2.
9 and 10	- Jobs A1 and A2 from client <i>DOG</i> finish on Drive 1. - However, jobs B1 and B2 for client <i>FOX</i> continue to run. - For Schedule A, the Media Multiplexing limit of 2 prevents job A5 from starting on Drive 1
11 and 12	- Job A3 from client <i>CAT</i> finishes on Drive 2. - Job B1 from client <i>FOX</i> finishes on Drive 1. - Job B2 is the only job currently running on Drive 1.
13	- Job A5 from client <i>CAT</i> starts on Drive 1. - Job A5 is the last job for Schedule A. - For Schedule A, the Media Multiplexing limit of 2 prevents job A5 from starting on Drive 2. - Therefore, job A5 starts on Drive 1.

NetBackup attempts to add multiplexed jobs to drives that already use multiplexing. If multiplexed jobs are confined to specific drives, other drives are available for non-multiplexed jobs.

If the backup window closes before NetBackup can start all the jobs in a multiplexing set, NetBackup completes only the jobs that have started.

For example, [Figure 20-13](#) assumes that the **Activity Monitor** shows jobs A1 through A5 as queued and active.

If only jobs A1 and A2 start before the window closes, NetBackup does not perform the other jobs that are in the set. If the window closes before any jobs start, only the first queued and active job starts and completes. In this example: Job A1.

About demultiplexing

Demultiplexing speeds up future restores and is useful for creating a copy for off-site storage. Use the duplication process in the **Catalog** utility to demultiplex a backup.

Duplication allows one multiplexed backup at one time to be copied from the source media to the target media. When duplication is complete, the target contains a single demultiplexed copy of each duplicated backup. (The target can also contain other backups.) The duplicate copy can be made into the primary copy. Do not

select **Preserve Multiplexing** in the **Configure Multiple Copies** dialog box when backups are duplicated.

Note: If you use the `bpduplicate` command instead of the **NetBackup Administration Console**, do not include the `-mpx` option on that command.

See [“Duplicating backup images”](#) on page 928.

Start Window tab

The **Start Window** tab provides controls for setting time periods during which NetBackup can start jobs when using a schedule. Time periods are referred to as windows. Configure windows so that they satisfy the requirements necessary to complete a job.

For example, create different windows:

- One for the backups that open each day for a specific amount of time
- Another for the backups that keep the window open all week

Adding, changing, or deleting a time window in a policy schedule

Use one of the following procedures to add, change, or delete a time window.

To add or change a time window in the NetBackup Administration Console

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Expand the policy name in the middle pane, and then select **Schedules**.
- 3 Do one of the following:

To add a time window	■ Click Actions > New > Schedule. ■ In the Add Schedule dialog box, enter the name of a schedule.

To change a time window	In the right pane, double-click the schedule you want to change. The Change Schedule dialog box appears.

- 4 Click the **Start Window** tab.


- 5 To change the increments available for selecting start times or end times, change the value in the **Resolution** field. You can choose 5, 10, 15, or 30 minutes. For example, a resolution of 10 minutes allows time window adjustments by 10-minute increments.
- 6 To indicate the opening of the time window, do the following:

Drag your cursor in the time table

Click the day and time when you'd like the window to start and drag it to the day and time when you'd like the window to close.

Use the settings in the dialog box

- In the **Start day** field, select the first day that the window opens.
- In the **Start time** field, select the time that the window opens.



- 7 To indicate the closing of the time window, do one of the following:

Drag your cursor in the time table

Click the day and time when you'd like the window to start and drag it to the day and time when you'd like the window to close.

Enter the duration of the time window Enter a length of time in the **Duration (days, hours, minutes)** fields.

Indicate the end of the time window

- Select a day in the **End day** list.
- Select a time in the **End time** field.

Time windows show as bars in the schedule display.

Specify enough time to allow all clients in the policy to complete a backup.

Consider allowing extra time in the schedule in case the schedule starts late due to factors outside of NetBackup. (Delays due to unavailable devices, for example.) Otherwise, all backups may not have a chance to start.

8 As necessary, do any of the following:

Click Delete .	Deletes the selected time window.
Click Clear .	Deletes all time windows from the schedule display.
Click Duplicate .	Replicates the time window for the entire week.
Click Undo .	Erases the last action.

9 Do one of the following:

Click Add .	To save the time window and leave the dialog box open.
Click OK .	To save the time window and close the dialog box.

Example of schedule duration

[Figure 20-14](#) illustrates the effect of schedule duration on two full backup schedules. The start time for Schedule B begins shortly after the end time for the previous Schedule A. Both schedules have three clients with backups due.

Figure 20-14 Duration example




Figure 20-14 illustrates the following points:

- Point 1 Client A3 starts within the Schedule A time window but doesn't complete until after the Schedule B start time. However, Client A3 runs to completion even if the window closes while the backup is running. Client B1, on Schedule B, begins as soon as Client A3 completes.
- Point2 Schedule A does not leave enough time for all the clients on Schedule B to be backed up. Consequently, Client B3 is unable to start because the time window has closed. Client B3 must wait until the next time NetBackup runs Schedule B.

Excluding days from a schedule

Use the **Exclude Days** tab to exclude specific days from a schedule for a backup policy. If a day is excluded from a schedule, jobs do not run on that day. The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.


To exclude a day from a schedule

- 1 Use one or more methods to indicate the days to exclude:
 - Select the day(s) on the 3-month calendar that you want to exclude. Use the drop-down lists at the top of the calendar to change the months or year.
 - To indicate **Recurring Week Days**:
 - Click **Select All** to select all of the days in every month for every year.
 - Click **Deselect All** to remove all existing selections.

- Check a box in the matrix to select a specific day to exclude for every month.
 - Click the column head of a day of the week to exclude that day every month.
 - Click the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to exclude that week every month.
 - To indicate **Recurring Days of the Month**:
 - Click **Select All** to select all of the days in every month.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select that day to exclude each month.
 - Click **Last Day** to exclude the last day of every month.
 - To indicate **Specific Dates**:
 - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.
The date appears in the **Specific Dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 2 Add additional dates as necessary, and then click **OK** to save the changes.

Include Dates tab

The **Include Dates** tab appears in the **Add New Schedule** or **Change Schedule** dialog box. For the tab to display, you must select the **Calendar** option as the **Schedule type** on the **Attributes** tab. Calendar-based schedules provide several run day options for determining when a task runs.

Figure 20-15 Calendar selection in the **Policy Attributes** tab

The **Calendar** attribute enables the **Include Days** tab

The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

Calendar scheduling with the Include Dates tab

Use the **Calendar** option in the policy **Attributes** tab to create a job schedule based on a calendar view. The **Include Dates** tab lets administrators configure the schedules that run according to specific days, on recurring week days, or on recurring days of the month.

Note: Using the calendar schedule, if a green checkmark does not appear on a day, the day is not included in the schedule.

If **Retries allowed after runday** is enabled, a job could run on a day that is not included in the schedule.

When a new calendar schedule is created with **Retries allowed after runday** enabled, the schedule runs its first job on the next day when the backup window is open. That day may be before the first run day that is included in the schedule.

To use a calendar to schedule run days

- 1 In the **Attributes** tab, enable the **Calendar** attribute.
- 2 Select the **Include Dates** tab.
- 3 Use one or more methods to schedule the days on which jobs can run:

- Select the day(s) on the three-month calendar that you want jobs to run. Use the drop-down lists at the top of the calendar to change the months or year.
 - To indicate **Recurring Week Days**:
 - Click **Select All** to select all of the days in every month for every year.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select a specific day to include for every month.
 - Click the column head of a day of the week to include that day every month.
 - Click the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to include that week every month.
 - To indicate **Recurring Days of the Month**:
 - Click **Select All** to select all of the days in every month.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select that day to include each month.
 - Click **Last Day** to include the last day of every month.
 - To indicate **Specific Dates**:
 - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.
 The date appears in the **Specific Dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 4 Add additional dates as necessary, and then click **OK** to save the included days.

How NetBackup determines which schedule to run next


When a policy contains one schedule, the schedule that is selected to run next is straightforward. But when a policy contains multiple schedules, choosing which schedule to run next can become more complicated.

NetBackup performs the following tasks to determine which schedule to run next:

- NetBackup determines the due time for each schedule. The due time depends on the following:

- The last backup data for each schedule based on comparable schedules.
- The frequency that is added to each schedule to determine which schedule is due next.
- NetBackup checks the start time for each schedule. The schedule with the soonest start time runs next. That is, the schedule with the next open window.

Figure 20-16 Schedule selection overview



When any of the following events occurs, NetBackup recalculates which schedule to run next in a policy:

- A backup job finishes.
- A client backup image expires.
- The Policy Execution Manager (`nbpem`) starts.
- The administrator changes the policy.
NetBackup looks for updated policies every 10 minutes. If the policy has recently been updated, NetBackup waits an additional minute to be sure that changes are not currently underway. You can change the frequency that NetBackup looks for updates by changing the **Policy Update Interval** in the **Global Attributes** host properties.

See [“Global attributes properties”](#) on page 109.

The due time for each schedule equals the last backup data for the schedule, plus the schedule’s frequency:

$$\text{Due time} = \text{Last backup data} + \text{Frequency}$$


Last backup data refers to the schedule that ran most recently among comparable schedules. NetBackup uses the date and time of that schedule to determine the due time for all the schedules that use that schedule as the last backup data.

In some cases, the last backup data for a schedule names the schedule itself. In other cases, the last backup data for a schedule is another comparable schedule. NetBackup makes the following comparisons to identify a comparable schedule:

Full schedules	Compared to other full schedules of the same or longer frequency.
Cumulative incremental schedules	Compared to the following: ■ Full schedules of the same or longer frequency. ■ Other cumulative incremental schedules of the same or longer frequency.
Differential incremental schedules	Compared to the following: ■ Full schedules of the same or longer frequency. ■ Cumulative incremental schedules of the same or longer frequency. ■ Other differential incremental schedules of the same or longer frequency. Note: To have a longer frequency means that the schedule is configured to run less often.

The comparison rules ensure that no schedule is overlooked for consideration, potentially causing a gap in backup coverage.

Figure 20-17 Schedule coverage



The following jobs create additional complexities in scheduling:

- Multistreaming jobs
- Each stream is scheduled independently. The data may change in the time between the streamed backups. Two restores that are based on the same backup may not be identical if created from different streams.
- Synthetic backup jobs
- NetBackup uses the previous synthetic job as the basis for determining when the next synthetic job should run.


About schedule windows that span midnight

A backup window may begin in one day and end in another. If this kind of policy is scheduled to run each day, NetBackup does not run the job again immediately after midnight. Instead, even though the window spans into the next day, NetBackup considers it to be one window. NetBackup recognizes that the administrator's intention is usually not to have a job run again so soon after the previous backup.

Figure 20-18 shows a window that spans midnight.

If a policy is scheduled to run each day, NetBackup looks to see if another window opens later in the day. If another window is set up to open later, NetBackup waits and runs the job then.

Figure 20-18 Schedule that spans midnight



The job is due Monday as well. Instead of running the job again immediately after midnight, NetBackup looks for a window later in the day and runs the job.


If no other window is scheduled to open later in the day, NetBackup does not wait. If the job has a daily frequency, the job runs again after midnight to meet the daily backup frequency requirement.

How open schedules affect calendar-based and frequency-based schedules

A single backup window can span the entire week. This kind of schedule is called an open schedule because a job may run at any time of day or night during the week. Open schedules affect calendar-based and frequency-based schedules differently.

[Figure 20-19](#) shows an open schedule.

Figure 20-19 An open schedule



Open schedules affect calendar-based and frequency-based schedules differently:

- Calendar-based schedules
- Calendar-based schedules run whenever the calendar schedule indicates. NetBackup assumes that an environment requires one backup on each day that is selected on the calendar schedule. Given an open schedule, backups run as soon after midnight as possible to satisfy the daily backup requirement.
- Frequency-based schedules
- Frequency-based schedules run when the frequency setting indicates. For example, with a frequency of one day, NetBackup runs backups at 24-hour intervals based on the start time.

Figure 20-20 shows that the backups on a calendar-based schedule would run Monday through Friday.

Figure 20-20 An open schedule that is calendar-based






Figure 20-21 and Figure 20-22 show that the backups based on a frequency-based schedule should run every day of the week, including Saturday and Sunday.

Figure 20-21 An open schedule that is frequency-based



In [Figure 20-22](#), backups run at 10:00 P.M. nightly based on the start time.

Figure 20-22 Example of a frequency-based schedule with an open schedule



Click on a window to see the start time and end time of each day.

The start time indicates when backups can run.

The window has a duration of 1 day. The window is duplicated for each day, to create an open schedule.

Creating an open schedule in the NetBackup Administration Console

The following procedure describes how to create an open schedule in an existing policy. In this procedure, the open schedule is configured to begin at 10:00 P.M.

To create an open schedule in the NetBackup Administration Console

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the middle pane, double-click on the policy name where you want to create an open schedule.
- 3 Select the **Schedules** tab.
- 4 Click **New** or **Add** to create a new schedule.
- 5 Complete the information on the **Attributes** tab.
- 6 Select the **Start Window** tab.
- 7 Select Sunday as the **Modify Day** and **10:00:00 PM** as the **Start time**.

- 8 Select Monday as the **End Day** and **10:00:00 PM** as the **End time**. The **Duration** is then automatically set to one day.
- 9 Click **Duplicate** to copy this window to each day of the week.
- 10 Click **OK** to add the schedule to the policy.

Runtime considerations that affect backup frequency

The following items may cause a NetBackup job to run more frequently than expected, or may prevent a job from meeting its backup frequency requirement.

Table 20-46 Items that can affect backup frequency

Item	Description
Changing a policy causes the policy to run	If the administrator changes or activates a policy, the change prompts NetBackup to run the job as soon as possible. It does not matter if the schedule is calendar-based or frequency-based.
Window availability	Whether the schedule is calendar-based or frequency-based, a job cannot run if windows are not open on the configured rundays. ■ For calendar-based schedules, windows must be open on the specific dates, recurring weekdays, or recurring days of the month that the calendar schedule indicates. Note: A frequency is not configurable for a calendar-based schedule. For this schedule type, NetBackup assumes a daily backup frequency. ■ For frequency-based schedules, a daily frequency requires that a window is open each day.
Backup attempt limit	A Global Attribute host property setting determines how many times a failed job can attempt to run. The Schedule backup attempts property includes the number of attempts and the time period in which the attempts can take place. By default, a failed job tries to run two times every 12 hours if an open window is available. Note that this setting supersedes any other frequency requirement and can cause a schedule to skip an open window. For example, if a job meets the maximum number of job attempts, NetBackup does not try to run the job again during the retry period indicated. It does not attempt, even in an open window and a daily backup frequency has not been met that day. See "Global attributes properties" on page 109.

About the Clients tab

The **Clients** tab contains a list of clients to be backed up (or acted upon) by the selected policy. A client must be included in the list of at least one backup policy to be backed up.

Placing a client in more than one backup policy can be useful. For example, place the client name in two policies to back up different sets of files on the client according to different policy rules.

The **Clients** tab does not appear for Vault or Catalog policy types.

Adding, changing, or deleting clients in a policy

A client must be included in the list of at least one active backup policy to be backed up. Use the following procedures to add, change, or delete clients in an existing NetBackup policy.

To add, change, or delete a client in a policy


- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Policies**.
- 2 Open the policy that you want to change.
- 3 Select the **Clients** tab and perform one of the following actions:

To add a new client	Continue to step 4.
To change an existing client	Double-click on the client that you want to change or select the client and click Change . The Change Client dialog box appears. Click OK to accept the change and close the Change Client dialog box. Or, click Cancel to close the dialog box without saving the changes.
To delete a client	Select a client and click Delete . Hold down Shift to select multiple clients. A confirmation dialog box appears that lists the clients to be deleted. Click Yes to delete the clients or No to escape the dialog box.

- 4 To add a new client, click **New**. The **Add Client** dialog box appears.

In the **Add Client** dialog box, enter the host name that you want to add. Or click the **Browse for Computer** button, select the host from the network tree, and click **OK**.

Note: The **Browse for clients** and the **Detect client operating system** options are unavailable for a BigData policy type. To add a client using the BigData policy, enter the name of the client, manually.



Observe the following rules for assigning client names:

- Use a name by which the server knows the client (one that you can use on the server to `ping` or `telnet` to the client).

Note: To add a client to backup universal share data, enter the host name of the client where the universal share is mounted. This name is used for cataloging. Although you can enter any name, a best practice is to enter the host short name, Fully Qualified Domain Name (FQDN), or IP address of the host that has permission to mount the universal share. For database systems, you can enter the host's network ID.

Note: To add a client for a Nutanix Acropolis Cluster, enter the display name of the virtual machine. The display name of a virtual machine is case-sensitive.

To add a client for a Hadoop cluster, enter the Fully Qualified Domain Name (FQDN) of the Hadoop cluster or namenode.

- You cannot add an identically named client twice to one policy. However, if you create a policy with **client_1** and **Client_1** as client names, NetBackup lets you save the policy. Update the `bp.conf` file using the `bpsetconfig` command. Set `CASE_INSENSITIVE_HOSTNAME_VALIDATION = YES` to force NetBackup to check for different character case in client names. The check is performed

before NetBackup saves the policy. The policy is not saved until the client name issue is fixed.

- If the client is in multiple policies, use the same name in each policy.
- If the network configuration has multiple domains, use a more qualified name. For example, use `client1.null.com` or `client1.null` rather than only `client1`.
- Add only clients with the hardware and the operating systems that this policy supports. For example, do not add a UNIX client to an **MS-Windows** policy. If you add a client to more than one policy, designate the same hardware and operating system in each of the policies.
 If the hardware and the operating system you want is not in the list, associated client software is not installed on the server. Check the `/usr/opensv/netbackup/client` directory for the directories and software that corresponds to the client you want to install. If the directories or software are not there, rerun the installation script on the server and select the option to install client software.
- To add a client to backup universal share data, enter the host name of the client where the universal share is mounted.
- Do not use an IP address as a client name in a policy or the backup may fail. Specify a host name instead.

- 5 You can either select the **Detect client operating system** check box or select the appropriate hardware and operating system for the client in the drop-down menu.
- 6 Perform one of the following:
 - Click **Add** to add the client to the list and leave the dialog box open to add another client.
 - Click **OK** to add the client to the list and close the dialog box.
 - Click **Cancel** to close the dialog box without adding the client.
- 7 When you are finished in the **Clients** tab:
 - Click **OK** to close and save the policy.
 - Click **Cancel** to close the policy without saving any additions or changes.

Browse for Hyper-V virtual machines

To browse for Hyper-V virtual machines

- 1
- Open the NetBackup web UI.
- 2
- Open the policy to which you want to add virtual machines.
- 3
- On the **Clients** tab, enter the **Server name**.
- 4
- Select **Select manually**.
- 5
- Select the **Add** button.

The following table describes the options that you can use to select Hyper-V virtual machines.

Table 20-47 Options for selecting Hyper-V virtual machines

Option	Description
Enter the VM display name	Note: The type of name to enter depends on the Primary VM identifier setting on the Hyper-V tab of the policy. Enter the host name, display name, or GUID of the virtual machine. The format of the host name or display name depends on your system. It may be the fully qualified name or another name, depending on your network configuration and how the name is defined in the guest OS. If NetBackup cannot find the name or GUID you enter, the policy validation fails. Make sure the Browse virtual machines option is cleared.
Browse virtual machines	Click this option to discover Hyper-V servers or cluster nodes (shown in the left pane). You can select virtual machines from a list (in the right pane). The virtual machine names that are listed may be derived from a cache file. Use of the cache file is faster than rediscovering the virtual machines on the network if your site has a large number of virtual machines. If the virtual machine is turned off but was turned on when the cache file was last created, its name appears in the list. If the display name of the virtual machine was recently changed in the Hyper-V Manager, note: The virtual machine name that was used for the backup does not change. If NetBackup cannot obtain the IP address of the virtual machine, the IP address is displayed as NONE.

Backup Selections tab

The **Backup Selections** tab contains a list of what to back up on each client, host, or instance when NetBackup runs an automatic schedule (for example, a full

backup). The list does not apply to user backups or archives, where users select the objects to back up before they start the operation.

The policy type determines what type of backup selections appear in the list. The backup selections list can contain the following:

- Paths that identify the location of files or directories
- Directives, which signal NetBackup to perform specific, predefined actions when it processes the selections list
- Scripts
See [“Registering authorized locations used by a NetBackup database script-based policy”](#) on page 789.
- Database objects
- Wildcards
Certain wildcards can be used in the selections list.
See [“Wildcard use in NetBackup”](#) on page 1055.
Windows clients support only the asterisk (*) and the question mark (?) as valid wildcards in the **Backup Selections** tab.
For information on how to use wildcards when you configure backup selections for database agents and other options, refer to the guide for that agent or option.

How the backup selections list is applied

The list applies to each client (or host, instance, or database) in the policy. Every file on the list does not need to exist on all of the clients. NetBackup backs up the files that it finds that are on the backup selections list. However, each client must contain at least one of the files in the backup selections list. Otherwise, the backup fails with status code 71 (none of the files in the file list exist).

Adding backup selections to a policy

Use the following procedures to add backup selections to a NetBackup policy, without opening up the tab view of the policy.

To add backup selections to a policy

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Policies**.
- 2 Open the policy where you want to add a backup selection.
- 3 Select the **Backup Selections** tab and perform one of the following actions:

Entering a path to a directory Click **Browse** to browse to a specific client to specify the path to the directory that you want to backup up. Or enter the pathname directly in the **Pathname or directive** field.

The **Browse** button toggles to an **Add** button after a selection is made. Click **Add** to add the path to the selection list.

A path may contain up to 1023 characters.

See [“Pathname rules for Windows client backups”](#) on page 795.

See [“Pathname rules for UNIX client backups”](#) on page 802.

Selecting a directive set or directive

Select or enter a directive or a directive set in the **Pathname or Directive set** drop-down menu.

Click **Add** to add the directive to the selection list.

See [“About the directives on the Backup Selections list”](#) on page 810.

Selecting a script

- Select or enter a script in the **Script** drop-down menu.

Scripts require that you specify the full path. Be sure that the scripts that are listed are installed on each of the clients that are specified on the **Clients** tab.

- For Oracle policies or DB2 policies, use the **Browse** button to locate the script for the **Script** list, or enter the name of the script.

Example of an Oracle script on UNIX:

```
install_path/netbackup/ext/db_ext/oracle/samples/  
rman/cold_database_backup.sh
```

Example of a DB2 script on UNIX:

```
/myscripts/db2_backup.sh
```

Click **Add** to add a script to the selection list.

You can select multiple directories and files for backup. The policy type determines whether the backup selections list can contain paths, directives, scripts, or a combination.

4 Click **OK** to add the selection list to the **Backup Selections** tab in the policy.

5 When you are finished in the **Backup Selections** tab:

- Click **OK** to close and save the policy.
- Click **Cancel** to close the policy without saving any additions or changes.

Creating a protection point for a NetBackup Appliance universal share

You can create a protection point for the data in a universal share that lets you manage and protect the data in the share. Creating a protection point is accomplished by creating a Universal-Share backup policy.

If a NetBackup Appliance is configured with multiple universal shares, a single policy can be created for some or all of the shares. You can also create individual policies, one for each share. If multiple appliances are configured with universal shares, each appliance should be configured with its own specific policy to protect the universal shares on that appliance.

For example, on an appliance, the NFS exported path (Universal Share) is `/shares/EXPORTED`. On the NFS client server, the mounted path of the Universal Share is `/mounted/MOUNTED`.

Where `/shares/EXPORTED` is the network path of the Universal Share data and `/mounted/MOUNTED` is the network path on the NFS client where the share is mounted.

You then back up the data on the `/mounted/MOUNTED` Universal Share.

To create a protection point policy for a NetBackup Appliance universal share

- 1 Create a Universal Share on the appliance using the NetBackup Appliance Web Console, NetBackup Appliance Shell Menu, or the NetBackup web UI.

For more information about the universal share feature and the supported appliance versions, refer to the NetBackup Appliance documentation:

<http://www.veritas.com/docs/DOC5332>

- 2 Mount the exported path of the Universal Share on the NFS client server.
- 3 Copy your application data to the Universal Share.
- 4 In NetBackup, on the **Attributes** tab, create a **Universal-Share** policy.

For the **Policy Storage**, you must use the storage unit that hosts the universal share. You must create one if one does not exist.

If multiple storage servers are configured with universal shares, each of the storage servers should be configured with its own specific policy. This configuration ensures that the universal shares on that storage server are protected.

- 5 On the **Schedules** tab, select either **FULL** or **INCR**.

Note: **Accelerator** backups are not supported or necessary for universal shares.

- 6 On the **Clients** tab, enter the name of the NetBackup Appliance where the Universal Share resides.

Note: Enter the host name of the client where the universal share is mounted. This name is used for cataloging. Although you can enter any name, a best practice is to enter the host short name, Fully Qualified Domain Name (FQDN), or IP address of the host that has permission to mount the universal share. For database systems, you can enter the host's network ID.

- 7 Select the **Backup Selections** tab and perform the following actions in the order listed:
 - (Optional) Add the `NEW_STREAM` directive if you require multistream backup. See “[NEW_STREAM directive](#)” on page 818.
 - Add the mounted path on the NFS client server and the exported path of the Universal Share.
`BACKUP /`
 - Enter the `BACKUP X USING Y` directive in the **Backup Selections** tab.
 For example: `BACKUP /mounted/MOUNTED USING /shares/EXPORTED`
 If the Universal Share is mounted on a Windows system,
`C:\mounted\MOUNTED`, use the `/C:/mounted/MOUNTED` format for the `BACKUP` path.
 You can add multiple shares in a policy. If you want to group several shares into one backup job, use the `NEW_STREAM` directive.
 See “[NEW_STREAM directive](#)” on page 818.
- 8 Enter the host name of the client where the Universal Share is mounted.
- 9 Run the **Universal-Share** policy.

After the backups are created, you can manage the backups with NetBackup features, such as restore, duplication, Auto Image Replication, and others.

You can immediately access the backups with NetBackup Instant Access APIs.

For information about NetBackup APIs, see the following website:

<https://sort.veritas.com/documents>

Select NetBackup and then the version at the bottom of the page.

Creating a protection point for a universal share

You can create a protection point for the data in a universal share that lets you manage and protect the data in the share. Creating a protection point is accomplished by creating a Universal-Share backup policy.

If an MSDP storage server is configured with multiple universal shares, a single policy can be created for some or all of the shares. You can also create individual policies, one for each share. If multiple storage servers are configured with universal shares, each storage server should be configured with its own specific policy to protect the universal shares on that storage server.

More information is available:

See [“Overview of universal shares”](#) on page 575.

To create a protection point policy for a universal share

- 1 Create a universal share on an existing MSDP storage server.
For details, see *Create a universal share* in [NetBackup Web UI Administrator's Guide](#).
- 2 Mount the exported path of the universal share on the storage server.
The **Export path** is found on the details page of the universal share in the NetBackup web UI: click **Storage > Disk storage > Universal shares** and then select the universal share to view its details.
- 3 Copy your application data to the universal share.
- 4 Create a policy using with the NetBackup web UI.
- 5 On the **Attributes** tab, select **Universal-Share** from the **Policy type** list.

Note: **Accelerator** backups are not supported or necessary for universal shares.

- 6 Under **Destination**, select the storage unit from the **Policy storage** list.

See *Policy storage (policy attribute)* in *NetBackup Administrator's Guide Volume I* for more information about policy storage setting.

For the **Policy storage**, you must use the storage unit that hosts the universal share. You must create one if one does not exist.

If multiple storage servers are configured with universal shares, each of the storage servers should be configured with its own specific policy. This configuration ensures that the universal shares on that storage server are protected.

The storage unit for universal share policy must be in the same disk pool volume where the universal share is created.

Note: If the primary server or the MSDP storage server is running NetBackup 10.0.1 or later, the media server must also be 10.0.1 or later.

- 7 On the **Schedules** tab, create a schedule.
- 8 On the **Clients** tab, enter the name of the desired client.

Universal share is an agentless technology, so the client name that is specified is used only for cataloging purposes. You can enter a NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance media server application instance, or MSDP BYO server name or a host where universal share is mounted. The client name can be a short name, Fully Qualified Domain Name (FQDN), or IP address.

- 9 On the **Backup selections** tab, enter the path of the universal share.

You can find the export path from the Universal share details page NetBackup web UI: **Storage > Disk storage > Universal shares**. For example:

```
/mnt/vpfs_shares/3cc7/3cc77559-64f8-4ceb-be90-3e242b89f5e9
```

You can use the `NEW_STREAM` directive if you require multistream backups.

See “[NEW_STREAM directive](#)” on page 818.

You can also use the `BACKUP X USING Y` directive, which allows cataloging under a different directory than the universal share path. For example: `BACKUP`

```
/demo/database1 USING
```

```
/mnt/vpfs_shares/3cc7/3cc77559-64f8-4ceb-be90-3e242b89f5e9. In this example, the backup is cataloged under /demo/database1.
```

- 10 Run the **Universal-Share** policy.

After the backups are created, you can manage the backups with NetBackup features, such as restore, duplication, Auto Image Replication, and others.

You can instantly access backup copies from local LSU or cloud LSU with web UI or NetBackup Instant Access APIs.

For information about NetBackup APIs, see the following website:

<https://sort.veritas.com/documents>

Select NetBackup and then the version at the bottom of the page.

Registering authorized locations used by a NetBackup database script-based policy

During a backup, NetBackup checks for scripts in the default script location and any authorized locations. The default, authorized script location for UNIX is

```
usr/opencv/netbackup/ext/db_ext
```

 and for Windows is

```
install_path\netbackup\dbext. If the script is not in the default script location
```

or an authorized location, the policy job fails. You can move any script into the default script location or any additional authorized location and NetBackup recognizes the scripts. You need to update the policy with the script location if it has changed. An authorized location can be a directory and NetBackup recognizes any script within that directory. An authorized location can also be a full path to a script if an entire directory does need to be authorized.

If the default script location does not work for your environment, use the following procedure to enter one or more authorized locations for your scripts. Use

`nbsetconfig` to enter an authorized location where the scripts reside. You can also use `bpsetconfig`, however this command is only available on the primary or the media server.

Note: One recommendation is that scripts should not be world-writable. NetBackup does not allow scripts to run from network or remote locations. All scripts must be stored and run locally. Any script that is created and saved in the NetBackup `db_ext` (UNIX) or `dbext` (Windows) location needs to be protected during a NetBackup uninstall.

For more information about registering authorized locations and scripts, review the knowledge base article:

https://www.veritas.com/content/support/en_US/article.100039639

To add an authorized location

- 1 Open a command prompt on the client.
- 2 Use `nbsetconfig` to enter values for an authorized location. The client privileged user must run these commands.

The following examples are for paths you may configure for the Oracle agent. Use the path that is appropriate for your agent.

- On UNIX:

```
[root@client26 bin]# ./nbsetconfig
nbsetconfig>DB_SCRIPT_PATH = /Oracle/scripts
nbsetconfig>DB_SCRIPT_PATH = /db/Oracle/scripts/full_backup.sh
nbsetconfig>
<ctrl-D>
```

- On Windows:

```
C:\Program Files\Cohesity NetBackup\NetBackup\bin>nbsetconfig
nbsetconfig> DB_SCRIPT_PATH=c:\db_scripts
nbsetconfig> DB_SCRIPT_PATH=e:\oracle\fullbackup\full_rman.sh
nbsetconfig>
<ctrl-Z>
```

Note: Review the [NetBackup Command Reference Guide](#) for options, such as reading from a text file and remotely setting clients from a NetBackup server using `bpsetconfig`. If you have a text file with the script location or authorized locations listed, `nbsetconfig` or `bpsetconfig` can read from that text file. An entry of `DB_SCRIPT_PATH=none` does not allow any script to run on a client. The `none` entry is useful if an administrator wants to completely lock down a server from running scripts.

- 3 (Conditional) Perform these steps on any clustered database or agent node that can perform the backup.
- 4 (Conditional) Update any policy if the script location was changed to the default or authorized location.

Verifying the Backup Selections list

Verify the **Backup Selections** list to make sure that the file paths are correct for the clients in the policy.

Table 20-48 Steps to verify the Backup Selections list

Step	Action	Description
Step 1	Check the syntax for the directives and the file path rules.	Do the following: ■ If the list includes directives, verify that the syntax for the directives is correct. ■ Check all entries against the file path rules for the clients in the policy. See "Pathname rules for Windows client backups" on page 795. See "Pathname rules for Windows disk image (raw) backups" on page 797. See "Pathname rules for Windows registry backups" on page 799. See "Pathname rules for UNIX client backups" on page 802. Path rules for the NetBackup clients that are running separately-priced options are covered in the NetBackup guide for the product. (For example, Snapshot Client or NetBackup for MS-Exchange.)
Step 2	Check for warning messages.	Do the following: ■ Run a set of backups. ■ Check the Problems report or the All log entries report for warning messages. The All logs entries report is also available in the NetBackup web UI. The backup status code does not always indicate errors on the Backup Selection list. Because NetBackup does not require all paths in the Backup Selections list to be present on all clients, an error may not be especially helpful. See "About the Reports utility" on page 1042.

Table 20-48 Steps to verify the Backup Selections list (*continued*)

Step	Action	Description
Step 3	Create a File System Backup Coverage Report .	Run the <code>check_coverage</code> script to create a File System Backup Coverage Report. The script is located in <code>install_path\NetBackup\bin\goodies</code> (on Windows) or in <code>/usr/openv/netbackup/bin/goodies</code> (on UNIX). The script can reveal mistakes in the selections list that make it impossible for NetBackup to find the files. Mistakes in the selections list can result in files being skipped in the backup. On Windows: If a path is not found, NetBackup logs a trivial (TRV) message or a warning (WRN) message. However, the same job can end with a backup status code of 0 (successful). Usually, to report files missing from the backup selections list is not helpful because not all files are expected to be present on every client. However, check the logs or use the <code>check_coverage</code> script to ensure that files are not missed due to bad or missing backup selections list entries. See "Example log messages from the File System Backup Coverage Report (check_coverage)" on page 792.

Example log messages from the File System Backup Coverage Report (check_coverage)

The **File System Backup Coverage Report** is created by running the `check_coverage` script. For information on `check_coverage`, see the comments in the script.

On Windows: The following example shows the log message that appears when files expected to be on a client are not found.

Assume that the backup selections list contains the path `c:\worklist` that is not present on all clients. NetBackup backs up `C:\worklist` on the clients where it exists.

For other clients, the **Problems** report or the **All log entries** report shows a message similar to the following:

```
9/1/14 8:28:17 AM carrot freddie Info from client freddie: TRV
- object not found for file system backup: C:\worklist
```

This message occurs if `c:\worklist` is not the correct path name. For example, the directory name is `c:\worklists`, but `c:\worklist` was typed.

Note: If the paths seem correct and the message appears, ensure that no trailing spaces appear in the paths.

On UNIX: The following table shows examples of the log messages that appear when files expected to be on a client are not found.

Table 20-49 Example UNIX log messages from the File System Backup Coverage Report

Example	Description
Regular expressions or wildcards	Assume that the backup selections list contains a regular expression: <pre>/home1[0123456789]</pre> NetBackup backs up /home10 through /home19 if both exist. If they are not present, the Problems report or the All log entries report displays a message similar to the following: <pre>02/02/14 20:02:33 windows freddie from client freddie: TRV - Found no matching file system for /home1[0123456789]</pre>
Path not present on all clients or wrong path specified	Assume that the backup selections list contains a path named /worklist that is not present on all clients. NetBackup backs up /worklist on the clients where it exists. For other clients, the Problems report or the All log entries report displays a message similar to the following: <pre>02/02/14 21:46:56 carrot freddie from client freddie: TRV - cannot process path /worklist: No such file or directory. Skipping</pre> This message occurs if /worklist is not the correct path name. For example, the directory name is /worklists, but /worklist was typed. Note: If the paths seem correct and the message continues to appear, ensure that no trailing spaces appear in the paths.

Table 20-49 Example UNIX log messages from the File System Backup Coverage Report (*continued*)

Example	Description
Symbolic link	Assume the backup selections list names a symbolic link. NetBackup does not follow symbolic links and provides a message in the Problems report or the All log entries report: <pre>02/02/14 21:46:47 carrot freddie from client freddie: WRN- /src is only being backed up as a symbolic link</pre> Resolve the symbolic link if you do not intend to back up the symbolic link itself.

How to reduce backup time

A client can be added to multiple policies, to divide the client's files among the different backup selections lists. Multiple policies can reduce the backup time for that client because the files can be backed up in parallel.

Multiple clients can be backed up in parallel in the following situations:

- Multiple storage devices are available (or if the policies are multiplexed).
- **Maximum jobs per client** (in **Global Attributes** host properties) and the **Limit jobs per policy** policy attribute are set to allow it.
See ["Global attributes properties"](#) on page 109.
See ["Limit jobs per policy \(policy attribute\)"](#) on page 678.

Note: Understand disk and controller input and output limitations before configuring including a client in multiple policies. For example, if two file systems overload the client when backed up in parallel, place both file systems in the same policy. Schedule the file systems at different times or set **Maximum jobs per client** to 1.

Another method to reduce backup time is to select **Allow multiple data streams** for a policy, and then add `NEW_STREAMS` directives to the backup selections list.

For example:

```
NEW_STREAM
file_a
file_b
file_c
NEW_STREAM
file_d
```

```
file_e
file_f
```

The example produces two concurrent data streams. The first data string contains `file_a`, `file_b`, and `file_c`. The second data stream contains `file_d`, `file_e`, and `file_f`.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 697.

Note: For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can cause longer backup times. The disk heads must move back and forth between the tracks that contain files for the respective streams.

A directive instructs NetBackup to perform specific actions to process the files in the backup selections list.

Pathname rules for Windows client backups

To back up Windows clients, use the following conventions for entries in the backup selections list.

Table 20-50 Pathname rules for Windows client backups

Item	Description
Paths per line	Enter one path per line.
Colons and backslashes	Begin all paths with the drive letter followed by a colon (:) and a backslash (\). To specify an entire volume, append a backslash (\) to the entry to ensure that all data is protected on that volume: Correct entry: <code>c:\</code> Incorrect entry: <code>c:</code>
Case sensitivity	The drive letter and path are case-insensitive. The following example entries would successfully indicate the same directory: <pre>c:\Worklists\Admin\ C:\worklists\admin\ c:\WORKLISTS\Admin\ C:\Worklists\ADMIN\</pre> Note: If a path is listed in the Backup Selections tab more than once, the data is backed up more than once.

Table 20-50 Pathname rules for Windows client backups (*continued*)

Item	Description
Wildcards	Asterisks (*) and question marks (?) are the only wildcard characters allowed in the backup selection list for Windows clients. Square brackets and curly brackets are not valid for Windows clients and can cause backups to fail with a status 71. For Windows clients, wildcards function correctly only when they are placed at the end of the path, in the file or directory name. For example: <pre>C:\abc\xyz\r*.doc</pre> Wildcard characters do not work elsewhere in the path. For example, an asterisk functions as a literal character (not as a wildcard) in the following examples: <pre>C:*\xyz\myfile</pre> <pre>C:\abc*\myfile</pre> See “Wildcard use in NetBackup” on page 1055.
All local drives	To back up all local drives except for those that use removable media, specify the following: <pre>: \</pre> Or <pre>*: \</pre> Or <pre>ALL_LOCAL_DRIVES</pre> The following drives are not backed up: floppy disks, CD-ROMs, and any drives that are located on remote systems but mounted on a system through the network.
Use of mapped drives	Do not specify a local drive path that is mapped to a CIFS share using the Windows Map Network Drive option. This holds true for a policy that contains multiple clients as well. Do not specify paths that point to different CIFS shares. For example: <pre>P:</pre> <pre>Q:</pre> <pre>R:</pre>


Table 20-50 Pathname rules for Windows client backups (*continued*)

Item	Description
Use of UNC path(s)	If a backup policy contains multiple clients that specify a UNC path as a backup selection, the redundant backup copies are created of the same data from different clients. Consider including the host in a policy as a client to be backed up. For example: <pre>\\host_name\dir1</pre>
Omitted or excluded files	By default, NetBackup does not back up some files. See “Files that are excluded from backups by default” on page 823. Exclude specific files from backups by creating an exclusion list on the client. See “About excluding files from automatic backups” on page 825. The following backup selection list uses Windows conventions: <pre>c:\ d:\workfiles\ e:\Special\status c:\tests*.exe</pre>

Pathname rules for Windows disk image (raw) backups

On Windows clients, you can back up a logical disk drive as a disk image. That is, NetBackup backs up the entire logical drive on a bit-by-bit basis rather than by directories and files. Use the **Full backup** backup type to perform a disk image backup.

Figure 20-23 Disk image backups



To specify a disk image backup, add the logical name for the drive to the policy backup selection list. Disk images can be included in the same backup selection list with other backups. In the following sample backup selection list, the first entry (\\.\c:) creates a disk image backup of a logical drive C.

```
\\.\c:
d:\workfiles\
e:\Special\status
HKEY_LOCAL_MACHINE:\
```

To restore the backup, the user clicks **Select for restore > Restore from Normal backup**.

When the backups are listed, the disk image appears as a file with the same name that was specified in the backup selection list. For the previous example, the file name would show as follows:

```
\\.\c:
```

When you enter the destination to restore the file, use the following format:

\\.*drive*:

Where *drive* is the location where the partition is to be restored.

Consider the following when working with disk image backups:

Windows Open File Backup methods	NetBackup first attempts to use Windows Open File Backup methods. If that fails, NetBackup locks the logical drive, which ensures that no changes occur during the backup. If there are open files on the logical drive, a disk image backup is not performed.
Open files	Before a disk image is backed up or restored, all applications that have a file opened on the partition should be shut down. If the applications are not shut down, the operation fails. Examples of such applications are Windows Explorer or Norton AntiVirus.
Copy-On-Write snapshots	Ensure that no active COW (Copy-On-Write) snapshots are in progress. If there is an active COW snapshot, the snapshot process itself has a handle open to the volume.
Raw partitions	NetBackup does not support raw partition backups on unformatted partitions.
Paging file	If the volume is configured to contain a paging file (<code>pagefile.sys</code>), a raw partition backup of that volume may fail. In order for a raw partition backup of that volume to succeed, the volume may need to be reconfigured so as not to contain a paging file. The raw partition backup of the volume may work without reconfiguration if a snapshot can successfully be taken of that volume.

Pathname rules for Windows registry backups

The Windows registry can be backed up for disaster recover or individual HKEYs can be backed up. Consider the following items when configuring a Windows registry backup.

Disaster recovery	To ensure a successful recovery in case of a disk failure, always back up the entire registry. That is, back up the directory that contains the entire registry. On most Windows systems, this directory is located at: <pre>%systemroot%\system32\config</pre> Where %systemroot% is the directory where Windows is installed. Note: To recover the registry, do not include individual registry files or HKEY entries in the selection list that's used to back up the entire registry. If you use a NetBackup exclude list for a client, do not exclude any registry files from your backups. To restore the registry in the case of a disk failure, see the disaster recovery chapter in the NetBackup Troubleshooting Guide.
Individual HKEYs	Do not back up individual HKEYs for disaster recovery. You cannot perform a disaster recovery by restoring HKEYs. Do not include HKEY entries in the same policy backup selection list that is used to back up the entire registry. However, to restore individual keys within the registry, create a separate policy, and then specify the specific HKEYs in the backup selection list for that policy. The following is an example HKEY entry for a policy backup selection list: <pre>HKEY_LOCAL_MACHINE:\</pre> Backups and restores are slower than if the entire registry was backed up.

About hard links to files and directories

A hard link is a directory entry for a file. Every file can be considered to have at least one hard link. A hard link differs from a symbolic link in that a hard link is not a pointer to another file. A hard link is two-directory entry that point to the same inode number.

If the backup selection list includes hard-linked files, the data is backed up only once during a backup. NetBackup uses the first file name reference that is found in the directory structure. If a subsequent file name reference is found, it is backed up as a link to the name of the first file. Back up the link that means only one backup copy of the data is created, regardless of the number of hard links. Any hard link to the data works.

On most UNIX systems, only the root user can create a hard link to a directory. Some systems do not permit hard links, and many vendors recommend that these

links be avoided. NetBackup does not back up and restore hard-linked directories in the same manner as files.

Hard-linked files and hard-linked directories are different in the following ways:

- During a backup, if NetBackup encounters hard-linked directories, the directories are backed up once for each hard link.
- During a restore, NetBackup restores multiple copies of the hard-linked directory contents if the directories do not already exist on the disk. If the directories exist on disk, NetBackup restores the contents multiple times to the same disk location.

On NTFS volumes or on UNIX systems, each file can have multiple hard links. Therefore, a single file can appear in many directories (or even in the same directory with different names). A volume serial number (VSN) and a File Index indicate the actual, unique file on the volume. Collectively, the VSN and File Index are referred to as the file ID.


During a backup, if the backup selection list includes hard-linked files, the data is backed up only once. NetBackup uses the first file name reference that is found in the directory structure. If a subsequent file name reference is found, the reference is backed up as a link to the name of the first file. To back up subsequent references means that only one backup copy of the data is created, regardless of the number of multiple hard links.

If all hard-link references are restored, the hard-linked files continue to point to the same ID as the other files to which they are linked. However, if all the hard links are not restored, you can encounter anomalies as shown in the following examples.

Example 1: Restoring Link2 and Link3


Assume that three hard links point to the same data. During a backup of Link2 and Link3, Link2 is encountered first and backed up. Then Link3 is backed up as a link to Link2. The three files are all hard-linked to the same data.

Figure 20-24 Example of hard links to the same data




The original copies of Link2 and Link3 are backed up to tape, and then deleted. Only Link1 is left on the disk.

Figure 20-25 Example of hard links backed up to tape and disk



During a subsequent restore, Link2 and Link3 are restored. The restored files, however, do not point to the same file ID as Link1. Instead, they are assigned a new file ID or inode number and the data is written to a new place on the disk. The data in the new location is an exact copy of what is in Link1. The duplication occurs because the backup does not associate Link2 and L3 with Link1.

Figure 20-26 Example of restored hard links



Example 2: Restoring Link3

Assume that this time you attempt to restore only Link3. However, NetBackup cannot link Link3 to Link2 because Link2 does not exist. The restore can complete only if it can link to Link2. A secondary restore request to the NetBackup server automatically restores Link2, which contains the data. Link2 can now be successfully restored.

Pathname rules for UNIX client backups

To back up UNIX clients, use the following conventions for entries in the backup selections list.

Table 20-51 Pathname rules for UNIX client backups

Item	Description
Pathnames per line	Enter one pathname per line. NetBackup supports a maximum path length of 1023 characters for UNIX clients.
Forward slash	Begin all pathnames with a forward slash (/).

Table 20-51 Pathname rules for UNIX client backups (*continued*)

Item	Description
Wildcard characters	The following wildcard characters are allowed: <pre> * ? [] { } </pre> For UNIX clients, wildcards can appear anywhere in the path. See “Wildcard use in NetBackup” on page 1055.
Trailing spaces	If a backup selection list entry contains trailing spaces and a matching entry is not found, NetBackup deletes the spaces and checks again. If a match is not found, NetBackup skips the entry and logs a message in the Problems report or the All log entries report: <pre> TRV - cannot process path pathname: No such file or directory. Skipping TRV - Found no matching file system for pathname </pre>
Mount points	Pathnames that cross mount points or that the client mounts through NFS can affect the backup configuration. Read about the Follow NFS and Cross mount points attributes before you create a backup selection list. See “Follow NFS (policy attribute)” on page 682. See “Cross mount points (policy attribute)” on page 685.
Bootable tapes	NetBackup can back up operating system, kernel, and boot files. However, NetBackup cannot create bootable tapes. Consult your system documentation to create a bootable tape.
Omitted or excluded files	By default, NetBackup does not back up all files. See “Files that are excluded from backups by default” on page 823. Exclude specific files from backups by creating an exclusion list on the client. See “About excluding files from automatic backups” on page 825.
Busy file settings	The Busy file settings host properties for UNIX clients offers alternatives for handling busy and locked files. See “Busy file settings properties” on page 59.
Access Control Lists (ACLs)	NetBackup backs up Access Control Lists (ACLs), where supported. See the <i>NetBackup Enterprise Server and Server OS Software Compatibility List</i> at the following URL: https://www.veritas.com/support/en_US/article.100040093

Table 20-51 Pathname rules for UNIX client backups (*continued*)

Item	Description
Sun PC NetLink	NetBackup can back up and restore Sun PC NetLink files.
Extended attribute files and named data streams	NetBackup backs up Extended attribute files and named data streams, where supported. See the NetBackup Software Compatibility List. See “About backing up and restoring extended attribute files and named data streams” on page 807.
VxFS extent attributes	On Hewlett-Packard and Solaris SPARC platforms, NetBackup backs up VxFS extent attributes.
Symbolic links	NetBackup backs up the symbolic link object and does not attempt to follow the link to back up what it may point to. To achieve a backup of the target of the symbolic link, include that target in the file list. Restoring the symbolic link object restores only the object and not the data to which it may point. To restore the target data, select it from the backup image. See “About hard links to files and directories” on page 800. Note: If NetBackup restores a symbolic link as <code>root</code>, NetBackup changes the owner and group to the original owner and group. When NetBackup restores a symbolic link as a non-root user, the owner and group are set to the owner and the group of the person who performs the restore. Resetting the owner and group does not cause problems. When the UNIX system checks permissions, NetBackup uses the owner and group of the file to which the symbolic link points.
Directory junctions	NetBackup backs up the directory junction object and does not attempt to traverse into the directory to which it may point. To achieve a backup of the target of the directory junction, include that target in the file list. Restoring the directory junction link object restores only the object and not the data to which it may point. To restore the target data, select it from the backup image.

See [“About the Reports utility”](#) on page 1042.

UNIX raw partitions

Save a copy of the partition table before a raw partition backup is performed. Retain the copy for reference. To restore the raw partition, make sure that a device file exists. Also, the partition where the table is restored must be large enough or the results of the restore are unpredictable.

Consider the following items when creating UNIX raw partition backups.

File changes during the backup	Use raw partition backups only if you can ensure that the files have not changed in any way during the backup. Or, in the case of a database, if you can restore the database to a consistent state by using transaction log files.
Backup archives	Do not perform backup archives of raw partitions on any client. An archive backs up the raw partition, and then deletes the device file that is associated with the raw partition. The file system does not recover the space that the raw partition uses.
File systems	Before backing up file systems as raw partitions, unmount the file system. Unmounting the file system allows buffered changes to be written to the disk. Also, it prevents the possibility of any changes in the file system during the backup. Use the <code>bpstart_notify</code> and the <code>bpend_notify</code> scripts to unmount and remount the backed-up file systems.
Mount points	The Cross mount points policy attribute has no effect on raw partitions. If the root partition is backed up as a raw partition and contains mount points to other systems, the file systems are not backed up. The other file systems are not backed up, even with Cross mount points selected. See “Cross mount points (policy attribute)” on page 685. The same is true for the Follow NFS policy attribute. NFS file systems that are mounted in a raw partition are not backed up. Nor can you back up raw partitions from other computers by using NFS mounts to access the raw partitions. The devices are not accessible on other computers through NFS. See “Follow NFS (policy attribute)” on page 682.
Disk volume managers	Specify the logical partition names for any disks that disk volume managers manage. (For example, Veritas Volume Manager (VxVM).
FlashBackup policy	For clients in a FlashBackup policy, refer to the NetBackup NAS Administrator's Guide and the NetBackup Snapshot Manager for Data Center Administrator's Guide for the differences between Standard and FlashBackup policies.

Windows Server Failover Clustering (WSFC) environment	The use of FlashBackup in a Windows Server Failover Clustering (WSFC) environment is supported, with the following limitation: Raw partition restores can only be performed when the disk being restored is placed in extended maintenance mode or removed from the WSFC resource group. Note: Early versions of WSFC do not allow extended maintenance mode functionality. If the cluster does not support placing disks in extended maintenance mode, it is still possible to perform raw restores to an alternate, non-shared disk.

If there are no file systems to back up and the disks are used in raw mode, back up the disk partitions as raw partitions. For example, databases are sometimes used in raw mode. Use `bpstart_notify` and `bpend_notify` scripts to provide the necessary pre-processing and post-processing of databases when they are backed up as raw partitions.

You can also perform a raw partition backup of a disk partition that is used for file systems. A disadvantage of this method is that you must restore the entire partition to recover a single file (unless FlashBackup is in use). To avoid overwriting the entire partition, use the redirected restore feature to restore the raw partition to another raw partition of the same size. Then, copy individual files to the original file system.

Raw partition backups are also useful for backing up entire disks. Since the file system overhead is bypassed, a raw partition backup is usually faster. The size of the raw partition backup is the size of the entire disk, regardless of whether the entire disk is used.

To specify a UNIX raw partition in the policy backup selection list, enter the full path name of the device file.

For example, on a Solaris system, enter:

```
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

Note: Do not specify wildcards (such as `/dev/rsd*`) in pathnames for raw partition backups. Doing so can prevent the successful restore of entire devices if there is overlap between the memory partitions for different device files.

You can include raw partitions in the same backup selection list as other backups. For example:

```
/home
/usr
```

```
/etc  
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

Note: NetBackup does not distinguish between full and incremental backups when it backs up a raw partition. The entire partition is backed up in both cases.

Raw partition backups occur only if the absolute pathname in the backup selection list is a block or character special device file. You can specify either block or character special device files. Character special device files are often faster because character devices avoid the use of the buffer cache for accessed disk data. Test both a block and character special device file to ensure the optimum backup speed for your platform.

Ensure that you specify the actual block-device or character-device files. Sometimes these are links to the actual device files. If a link is specified, only the link is backed up. If the device files are reached while backing up `/dev`, NetBackup backs up only the inode files for the device, not the device itself.

To perform a raw partition backup, select `Full backup` for the **Type of backup** from the **Schedules** tab. Any other backup type does not work for backing up raw partitions.

See [“Type of backup \(schedule attribute\)”](#) on page 732.

About backing up and restoring extended attribute files and named data streams

NetBackup can back up and restore the following file attributes:

- Extended attribute files of the Solaris UNIX file system (UFS) and temporary file system (tmpfs)
- Named data streams of the VxFS file system

NetBackup backs up extended attribute files and named data streams as part of normal file system backups.

Extended attribute files and named data streams are normal files contained in a hidden attribute directory that relate to a particular base file. The hidden directory is stored within the file system, but can be accessed only by the base file to which it is related. To view which files have extended attributes on Solaris 9 (or greater) systems, enter: `ls -@`


Neither extended attribute files nor named data streams can be backed up or restored individually. Rather, the files are backed up and restored all at once along with the base file.

The presence of a large number of extended attribute files or named data streams can cause some degradation in backup and restore speed. The speed is affected since the base file and all associated files are backed up.

The speed is especially affected in the case of incremental backups, during which NetBackup checks the `mtime` or `ctime` of each file individually.

On UNIX:

Figure 20-27 Example of base file and extended attribute directory and files



To back up or restore named data streams and extended attributes, the client, media server, and primary server must run the following versions:

- NetBackup clients
 - HP 11.23 running VxFS 4.1 or greater.

Note: Access Control Lists (ACLs) are not backed up unless running VxFS 5.0 or greater.

- AIX running VxFS 4.0 or greater.

Note: Access Control Lists (ACLs) are not backed up unless running VxFS 5.0 or greater.


- Solaris 10 running VxFS 5.0 or greater
- Solaris SPARC 9, 10 running VxFS 4.0 or greater
- Linux running VxFS 5.0 or greater.
- A NetBackup primary server

A NetBackup primary server of any version can back up and restore named data streams and Solaris extended attributes.

Restored attribute files and named data streams can replace existing files if **Overwrite existing files** is selected in the **Backup, Archive, and Restore** client interface.


On UNIX: In the following example, File 1 is to be restored. Base File 1 currently possesses four extended attribute files.

Figure 20-28 Extended attribute files of Base File 1



On UNIX: The user restores File 1 from a backup that was created when File 1 possessed only three extended attribute files.

Figure 20-29 Backup of File 1



On UNIX: Since **Overwrite existing files** is selected as a restore option, when the user restores File 1, extended attribute files 1, 2, and 3 are overwritten. Extended attribute file 4 remains and is not overwritten.

Figure 20-30 Restore of File 1



If an attempt is made to restore the following items, an error message appears in the **Restore Monitor**. The error message informs the user that the extended attributes or named data streams are not restored.

- The extended attribute files to any non-Solaris 9 client (or greater)

- Named data streams to any non-VxFS 4.0 client

NetBackup then continues with the restore job.

To disable the restore of extended attribute files and named data streams, add an empty file to the client. Name the file `IGNORE_XATTR` and place it in the following directory:

```
/usr/opensv/netbackup/
```

The addition affects only Solaris 9 or VxFS 4.0 clients.

File `IGNORE_XATTR` was formerly known as `IGNORE_XATTR_SOLARIS`.

On UNIX: Only `nbtar` that is supplied with NetBackup can restore the extended attributes or named data streams to a client.

On UNIX: For more information, see the [NetBackup Administrator's Guide, Volume II](#).

Note: Extended attributes and named data streams cannot be compressed.

Pathname rules for the clients that run extension products

Path rules for the NetBackup clients that are running separately-priced options are covered in the NetBackup guide for the product. (For example, Snapshot Client or NetBackup for MS-Exchange.)

About the directives on the Backup Selections list

Directives on the **Backup Selections** list signal NetBackup to perform specific, predefined actions when it processes the files on the selections list.

The available directives depend on the policy type and whether the **Allow multiple data streams** attribute is enabled for the policy. The following example is a backup selections list that contains the `NEW_STREAM` directive. The **MS-Windows** policy type is selected, and **Allow multiple data streams** is enabled.

```
NEW_STREAM
D:\Program Files
NEW_STREAM
C:\Winnt
```

Note: For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. The heads must move back and forth between the tracks that contain files for the respective streams.

The following table summarizes many of the directives available on the **Backup Selections** list.

Table 20-52 Summary of directives on the Backup Selections list

Directive	Description	Applicable operating system
ALL_LOCAL_DRIVES	Instructs NetBackup to back up all local drives except for those drives that use removable media. This directive also includes critical system-related components. See “ ALL_LOCAL_DRIVES directive ” on page 814.	All supported systems
System State:\	Instructs NetBackup to back up critical system-related components. The exact set of system components that are backed up depends on the operating system version and system configuration. See “ System State:\ directive ” on page 816.	All supported Windows systems
Shadow Copy Components:\	Instructs NetBackup to back up all writers for the Volume Shadow Copy component. This also implies and/or includes System State:\ if that was not also selected. See “ Shadow Copy Components:\ directive ” on page 816.	All supported Windows systems
Active Directory Application Mode:\	Active Directory Application Mode (ADAM) is a lightweight directory service that runs as a user service. This directive can be used to back up ADAM data on computers where it is installed. However, it does not back up the Active Directory itself.	All supported Windows systems
Policy-specific directives	Apply only to specific policy types and can appear only in backup selections lists for those policies. See “ Directives for specific policy types ” on page 817.	Policy type specific
UNSET and UNSET_ALL	Interrupt the streaming of policy-specific directives. The Allow multiple data streams policy attribute must be enabled before these directives can be used. See “ UNSET and UNSET_ALL directives ” on page 822.	All

Table 20-52 Summary of directives on the Backup Selections list (*continued*)

Directive	Description	Applicable operating system
NEW_STREAM	When NEW_STREAM is on the first line of the Backup Selections list, this directive determines how a backup is performed in the following modes: ■ Administrator-defined streaming ■ Auto-discovery streaming The Allow multiple data streams policy attribute must be enabled before this directive can be used. See “NEW_STREAM directive” on page 818.	All
USHARE	Instructs NetBackup to back up the Universal Share data. The Allow multiple data streams policy attribute must be enabled before this directive can be used. See “USHARE directive” on page 818.	

ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives

The **ALL_FILESYSTEMS** directive provides a method to include all file systems and volumes on an NDMP filer in an NDMP backup policy.

You can exclude specific volumes from an **ALL_FILESYSTEMS** backup selection if you do not want to back up every volume on an NDMP filer. Use the **VOLUME_EXCLUDE_LIST** directive for this purpose. You may use valid wildcard characters in the **VOLUME_EXCLUDE_LIST** statement.

Note: The following examples use selections that are specific to NetApp Data ONTAP 7-mode. For specific examples of backup selections for other configurations, refer to the appropriate documentation.

The **VOLUME_EXCLUDE_LIST** statements must precede **ALL_FILESYSTEMS** statement. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01
ALL_FILESYSTEMS
```

or


```
VOLUME_EXCLUDE_LIST=/vol/testvol*
ALL_FILESYSTEMS
```

To specify multiple values in a `VOLUME_EXCLUDE_LIST` statement, separate the values with a comma. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01,/vol/testvol*
ALL_FILESYSTEMS
```

You can also specify more than one `VOLUME_EXCLUDE_LIST` statement with an `ALL_FILESYSTEMS` directive. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01
VOLUME_EXCLUDE_LIST=/vol/testvol*
ALL_FILESYSTEMS
```

A `VOLUME_EXCLUDE_LIST` statement may include a maximum of 256 characters. Create multiple `VOLUME_EXCLUDE_LIST` statements if necessary to avoid exceeding the limit of 256 characters. If you specify more than 256 characters, the volume list is truncated. A truncated statement may result in a backup job failure, and the error message `Invalid command parameter(20)` is displayed.

If the backup selection includes read-only volumes or full volumes, an NDMP backup job fails with the status code 20 (`Invalid command parameter(20)`). If you encounter a similar NDMP backup job error, review the `ostfi` logs to identify the volumes for which the failure occurred. You can use `VOLUME_EXCLUDE_LIST` statements with the `ALL_FILESYSTEMS` statement to exclude the read-only volumes and the volumes with insufficient space.

In a NetBackup Replication Director environment where snapshots are replicated to a secondary filer, it is recommended that you use storage lifecycle policies to control backups on the secondary filer.

On NetApp 7-mode storage systems, it is generally not recommended for users to store files in `/vol/vol0` because the volume contains filer system files. For this reason, `vol0` should be excluded from the backup if the `ALL_FILESYSTEMS` directive is used in the backup policy. The following is a backup selection list that excludes `/vol/vol0`:

```
VOLUME_EXCLUDE_LIST=/vol/vol0
ALL_FILESYSTEMS
```

- Do not use `ALL_FILESYSTEMS` to backup all volumes on a secondary filer. Inconsistencies may occur when automatically created NetApp FlexClone volumes are backed up or restored. Such volumes are temporary and used as

virtual copies or pointers to actual volumes and as such do not need to be backed up.

- If you must back up all volumes on a secondary filer, it is recommended that you exclude the FlexClone volumes as well as replicated volumes. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Clone_*
VOLUME_EXCLUDE_LIST=/vol/*_[0-9]
VOLUME_EXCLUDE_LIST=/vol/*_[0-9][0-9]
VOLUME_EXCLUDE_LIST=/vol/*_[0-9][0-9][0-9]
ALL_FILESYSTEMS
```

This example assumes all FlexClone volumes and only FlexClone volumes begin with `/vol/Clone_`. Adjust the volume specifications appropriately for your environment.

- `VOLUME_EXCLUDE_LIST` applies only to `ALL_FILESYSTEMS`. It does not apply to explicit backup selections or wildcard-based backup selections. If you use the `ALL_FILESYSTEMS` directive in an NDMP policy for Clustered Data ONTAP, you must exclude each selected SVM's root volume using the `VOLUME_EXCLUDE_LIST` directive. Otherwise the backups fail.

Backups from snapshots for NDMP policies fail when the import of a snapshot fails for volumes where logical unit numbers (LUNs) reside with status code 4213 (Snapshot import failed). To avoid this error, use the `VOLUME_EXCLUDE_LIST` directive to exclude any volumes that are used to create LUNs accessed through a storage area network (SAN).

ALL_LOCAL_DRIVES directive

Use the `ALL_LOCAL_DRIVES` directive to back up all local drives except for those drives that use removable media. If this directive is used, this directive must be the only entry in the backup selections list for the policy. No other files or directives can be listed. The directive applies only to the following policy types:

- MS-Windows
- Standard

`ALL_LOCAL_DRIVES` gives different results depending on whether **Allow multiple data streams** is enabled for the policy:

Allow multiple data streams enabled	Applies only to Standard and MS-Windows policy types. NetBackup backs up the entire client, and then splits the data from each drive (Windows) or file system (UNIX) into its own backup stream. NetBackup periodically preprocesses the client to make necessary changes to the streams.

Allow multiple data streams NetBackup backs up the entire client and includes all drives disabled and file systems in the same stream.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 697.

Caution: Do not select **Cross mount points** for policies where you use the `ALL_LOCAL_DRIVES` directive.

See [“ALL_LOCAL_DRIVES example: Auto-discovery mode”](#) on page 815.

See [“ALL_LOCAL_DRIVES example: Without multiple data streams”](#) on page 815.

ALL_LOCAL_DRIVES example: Auto-discovery mode

Assume that **Allow multiple data streams** is enabled in the auto-discovery mode. Assume that the client is a Windows system with two drive volumes, `C:\` and `D:\`. The backup selections list contains the following directive:

```
ALL_LOCAL_DRIVES
```

For this backup selections list, NetBackup generates the following:

- One stream for `C:\`
- One stream for `D:\`

For a UNIX client, NetBackup generates a stream for each file system.

`SYSTEM_STATE` is also backed up because `SYSTEM_STATE` is included in the `ALL_LOCAL_DRIVES` directive.

See [“ALL_LOCAL_DRIVES example: Without multiple data streams”](#) on page 815.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 697.

ALL_LOCAL_DRIVES example: Without multiple data streams

Assume that **Allow multiple data streams** is not enabled. Assume that the client is a Windows system with two drive volumes, `C:\` and `D:\`. The backup selections list contains the following directive:

```
ALL_LOCAL_DRIVES
```

For this backup selections list, NetBackup backs up the entire client in one data stream that contains the data from both `C:\` and `D:\`.

`SYSTEM_STATE` is also backed up because `SYSTEM_STATE` is included in the `ALL_LOCAL_DRIVES` directive.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 697.

System State:\ directive

The `System State:\` can be used on all supported Windows systems.

The `System State:\` directive is needed for the operating system versions which do not support Shadow Copy Components.

The `System State:\` directive creates a backup for critical system-related components. The exact set of system components that are backed up depends on the operating system version and system configuration.

The list of items that are backed up can include the following:

- Active Directory
- COM+ Class Database
- Cluster Database
- IIS Database
- Registry
- Boot Files and protected files
- SYSVOL
- Certificate Server

The files that comprise the registry can be found in the following location:

```
%SystemRoot%\SYSTEM32\Config
```

At a minimum, the following files are backed up as part of the registry:

- DEFAULT
- SAM
- SOFTWARE
- SECURITY
- SYSTEM

Shadow Copy Components:\ directive

The `Shadow Copy Components:\` directive specifies that all of the Volume Shadow Copy component writers get backed up. This directive affects the backups of the following clients:

- Windows 2003 Server computers that use the Volume Shadow Copy components.
- Windows IA64 systems with EFI System partitions.

Note: In the policies that back up Windows clients on IA64 platforms, use the `Shadow Copy components:\` directive instead of the `System_State:\` directive. The `Shadow Copy components:\` directive includes System State components and the EFI System partition automatically in the backup.

Since the Shadow Copy Components contain System State information, the Shadow Copy Components need to be backed up by a full backup.

The Volume Shadow Copy components include the following:

System State writers	- System files - COM+ Class Registration Database - SYSVOL - Active Directory - Cluster quorum - Certificate Services - Registry - Internet Information Services
System Service writers	- Removable Storage Manager - Event logs - Windows Internet Name Service - Windows Management Instrumentation - Remote Storage - Dynamic Host Configuration Protocol - Terminal Server Licensing - Background Intelligent Transfer Service
User Data	Items that the computer does not require to operate. For example, Active Directory Application Mode and Microsoft Distributed File System Replication (DFS) folders. See “About Microsoft DFSR backups and restores” on page 654.
Other Data	A category that is intended for future NetBackup releases.

Directives for specific policy types

Some directives apply only to specific policy types and can appear only in backup selections lists for those policies. NetBackup passes policy-specific directives to the clients along with the backup selections list. The clients then perform the appropriate action according to the directive. All policy-specific directives that are passed to a client in a stream are passed in all subsequent streams.

Note: Include policy-specific directives only in backup selections lists for the policies that support the directives or errors can occur.

The following policy types have their own backup selections list directives:

- FlashBackup
- NDMP
- Lotus-Notes
- MS-Exchange-Server

For information on other policy types and associated backup selections list directives, see the NetBackup guide for the option.

USHARE directive

The `USHARE` directive is recognized only if **Allow multiple data streams** is set for the Standard policy. `USHARE` directives are ignored if **Allow multiple data streams** is not set. `USHARE` must be on the first line if `NEW_STREAM` is not added.

If `NEW_STREAM` is added, `USHARE` must be in second line. The presence of `USHARE` on the first line or second line of the backup selections list determines the backup type as `USHARE` backup.

NEW_STREAM directive

The `NEW_STREAM` directive is recognized only if **Allow multiple data streams** is set for the policy. `NEW_STREAM` directives are ignored if **Allow multiple data streams** is not set.

If this directive is used in a backup selections list, the first instance of it must be on the first line. If it appears on the first line, it can also appear elsewhere in the list.

The presence of `NEW_STREAM` on the first line of the backup selections list determines how the backup is performed in the following modes: in administrator-defined streaming or in the auto-discovery streaming.

About the administrator-defined streaming mode

If `NEW_STREAM` is the first line of the backup selections list, the backup is performed in the administrator-defined streaming mode.

In this mode, the following actions occur:

- The backup splits into a separate stream at each point in the backup selections list where the `NEW_STREAM` directive occurs.

- All file paths between `NEW_STREAM` directives belong to the same stream.
- The start of a new stream (a `NEW_STREAM` directive) defines the end of the previous stream.
- The last stream in the backup selections list is terminated by the end of the backup selections list.

In the following examples, assume that each stream is from a separate physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. The backup time is longer if the heads must move back and forth between the tracks that contain files for the respective streams.

For example, consider the following backup selections list:

On Windows:

```
NEW_STREAM
D:\Program Files
C:\Winnt
NEW_STREAM
C:\users
D:\DataFiles
```

This backup selections list contains two data streams:

- The `NEW_STREAM` directive at the top of the list starts administrator-defined streaming and the first data stream. This stream backs up `D:\Program Files` and `C:\Winnt`.
- The second `NEW_STREAM` starts a second data stream that backs up `C:\users` and `D:\DataFiles`.

On UNIX:

```
NEW_STREAM
/usr
/lib
NEW_STREAM
/home
/bin
```

This backup selection list contains two data streams:

- The `NEW_STREAM` directive at the top of the list starts administrator-defined streaming and the first stream. This stream backs up `/usr` and `/lib`.
- The second `NEW_STREAM` starts a second data stream that backs up `/home` and `/bin`.

If a backup selections list entry is added to a stream, the entry is not backed up until the schedule is due for the policy. If the next backup due is an incremental, only the files that changed are backed up. To ensure that a new entry gets a full backup the first time, add it to a new stream. NetBackup performs a full backup of new streams that are added to the backup selections list.

In the previous example, assume that you add the following:

On Windows:

```
D:\Utilities
```

On UNIX:

```
/var
```

after

On Windows:

```
D:\Datafiles
```

On UNIX:

```
/bin
```

If an incremental backup is due that night, only changed files in `D:\Utilities` (on Windows) or in `/var` (on UNIX) are backed up. Add a `NEW_STREAM` directive before `D:\Utilities` (on Windows) or before `/var` (on UNIX), to perform a full backup of all files in `D:\Utilities` (on Windows) or in `/var` (on UNIX), regardless of when the files were last changed.

About the auto-discovery streaming mode

The auto-discovery streaming mode is initiated if the `NEW_STREAM` directive is not the first line of the backup selections list. The list must contain either the `ALL_LOCAL_DRIVES` directive or wildcards.

In this mode, the backup selections list is sent to the client, which preprocesses the list and splits the backup into streams as follows:

- If the backup selections list contains the `ALL_LOCAL_DRIVES` directive, NetBackup backs up the entire client. However, NetBackup splits each drive volume (Windows) or file system (UNIX) into its own backup stream.
See [“ALL_LOCAL_DRIVES directive”](#) on page 814.
- If wildcards are used, the expansion of the wildcards results in one stream per wildcard expansion. Wildcard usage is the same as for Windows clients.
See [“Wildcard use in NetBackup”](#) on page 1055.

If the backup selections list contains neither the `ALL_LOCAL_DRIVES` directive nor wildcards, the auto-discovery mode is not used. The server preprocesses rather than the client. Each file path in the backup selections list becomes a separate stream.

The auto-discovery streaming mode applies to Standard and MS-Windows policy types.

Before the backup begins, the client uses auto-discovery to preprocess the backup selections list to determine how many streams are required. The first backup that a policy performs preprocesses the backup selections list. Depending on the length of the preprocess interval, preprocessing may not occur before every backup.

About setting the preprocess interval for auto-discovery

The preprocess interval applies only to auto-discovery mode and specifies how often preprocessing occurs. When a schedule is due and NetBackup uses auto-discovery, NetBackup checks whether the previous preprocessing session has occurred within the preprocess interval.

NetBackup performs one of the following actions:

- If the preprocessing session occurs within the preprocess interval, NetBackup does not run preprocessing on the client.
- If the preprocessing session did not occur within the preprocess interval, NetBackup preprocesses the client and makes required changes to the streams.

If necessary, you can change the interval by using the `bpconfig` command. The default is 4 hours and is a good value for most of the sites that run daily backups.

If the interval is too long or too short, the following problems can occur:

Interval is too long.	Can cause missed backups because new streams are not added early enough. For example, assume that the preprocess interval is set to 4 hours and a schedule has a frequency of less than 4 hours. A new stream can be omitted from the next backup because the preprocessing interval has not expired when the backup is due.
Interval is too short.	Can cause preprocessing to occur often enough to increase scheduling time to an unacceptable level. A short interval is most likely to be a problem when the server must contact a large number of clients for preprocessing.

Use the following form of the `bpconfig` command to change the interval:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpconfig [-prep hours]
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpconfig [-prep hours]
```

For more information on the `bpconfig` command, see the [NetBackup Commands Reference Guide](#).

UNSET and UNSET_ALL directives

UNSET, UNSET_ALL The `UNSET` and `UNSET_ALL` directives interrupt the streaming of policy-specific directives.

All policy-specific directives that are passed to a client in a stream are passed in all subsequent streams. The `UNSET` and `UNSET_ALL` directives change this behavior. These directives are recognized only if the **Allow multiple data streams** option is set for the policy.

See [“Directives for specific policy types”](#) on page 817.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 697.

UNSET

The `UNSET` directive interrupts a policy-specific directive so it is not passed with any additional streams. The directive that was unset can be defined again later in the backup selections list to be included in the current and the later streams.

In the following backup selections list, the `set` command is a client-specific directive that is passed to the first and all subsequent streams.

```
NEW_STREAM
set destpath=/etc/home
/tmp
/use
NEW_STREAM
/export
NEW_STREAM
/var
```

For the `set` command to be passed to the first two streams only, use `UNSET` or `UNSET_ALL` at the beginning of the third stream. At this location, it prevents `SET` from being passed to the last stream.

```
NEW_STREAM
set destpath=/etc/home
/tmp
/use
NEW_STREAM
/export
NEW_STREAM
UNSET set destpath=/etc/home [or UNSET_ALL]
/var
```

UNSET_ALL

`UNSET_ALL` has the same effect as `UNSET` but unsets all policy-specific directives in the backup selections list that have been defined up to this point.

Files that are excluded from backups by default

By default, a number of files and file states are not backed up by NetBackup.

You can also exclude specific files from automatic backups by specifying the files or directories in an exclude list on the client.

See [“About excluding files from automatic backups”](#) on page 825.

By default, NetBackup does not back up the following files:

- NFS files or directories. To back up NFS files, enable **Follow NFS**.

- Files or directories in a different file system. To back up files in a different file system, enable **Cross mount points**.
- Files or directories with path lengths longer than 1023 characters.
- Files or directories in which the operating system does not return inode information (the `lstat` system call fails).
- Directories that NetBackup cannot access (the `cd` command cannot access).
- Socket special files. (Named pipes are backed up, however.)
- Locked files when locked by an application that currently has the file open.
- Busy files. If a file is open, NetBackup backs up the last saved version of the file.
- Files or directories beneath a `bind` mount (Linux).

NetBackup automatically excludes the following file system types on most platforms:

- `cdrom` (all UNIX platforms)
- `cachefs` (AIX, Solaris, UnixWare)
- `devpts` (Linux)
- `mntfs` (Solaris)
- `proc` (UNIX platforms)
Does not exclude automatically for AIX, so `/proc` must be added manually to the exclude list. If `/proc` is not added manually, partially successful backups may result with the `ALL_LOCAL_DRIVES` directive on AIX.
- `tmpfs` (Linux)
- `usbdevfs` (Linux)

See [“Follow NFS \(policy attribute\)”](#) on page 682.

See [“Cross mount points \(policy attribute\)”](#) on page 685.

About host identity-specific files excluded from the backup

To proactively avoid vulnerabilities, certain host identity-specific files are excluded from the backup.

- To identify the files that are not backed up, you can run one the following commands:
 - `nbgetconfig -private_exld_list`
 - `bpgetconfig -private_exld_list`

For more information about the commands see the [NetBackup Commands Reference Guide](#).

- Including the files in the backup:
If you don't want to exclude certain files from the backup, you must include those files in the include list.
For more information, see See [“About excluding files from automatic backups”](#) on page 825.
- Recreating the files that were not backed up:
The files that are not backed up are not restored. Ideally, the files would reside on the original location. However, in case you want to recreate the files, some of the key files and certificates can be recreated by restarting the NetBackup services. If you encounter any error that is related to keys or certificates, restart the NetBackup services and verify if the key files or the certificate is recreated. If the key file or certificate is not created, proceed with the certificate and key regeneration procedures that are provided in the [NetBackup Commands Reference Guide](#).

About excluding files from automatic backups

On most NetBackup clients, you can exclude specific files from automatic backups by specifying the files in an exclude list on the client.

You can also create an include list to add a file(s) specifically that is excluded. The include list is useful when, for example, an entire directory is excluded except for one file on the include list.

Note: Exclude and include lists do not apply to user backups and archives.

The method for specifying files in the exclude list and the include list depends on the type of client as follows:

Microsoft Windows clients Specify exclude and include lists in the **Backup, Archive, and Restore** client interface. Start **Backup, Archive, and Restore**. On the **File** menu, click **NetBackup Client Properties**. Select the **Exclude List** tab or the **Include List** tab. For further instructions, see the NetBackup user's guide for the client.

The **Exclude List** or the **Include List** can also be specified through the **NetBackup Administration Console** on the primary server.

See [“Exclude list properties”](#) on page 94.

UNIX clients	Create the exclude and include lists in the following files on the client: ■ <code>/usr/opensv/netbackup/include_list</code> ■ <code>/usr/opensv/netbackup/exclude_list</code>
Specific policy	Create an exclude list for a specific policy or for a policy and a schedule combination. Create an <code>exclude_list</code> file with a <code>.policyname</code> or <code>.policyname.schedulename</code> suffix. The following two file examples use a policy that is named <code>wkstations</code>. The policy contains a schedule that is named <code>fulls</code>: <pre> /usr/opensv/netbackup/exclude_list.wkstations /usr/opensv/netbackup/exclude_list.wkstations.fulls </pre> The first file affects all scheduled backups in the policy that is named <code>wkstations</code>. The second file affects backups only when the schedule is named <code>fulls</code>. For a given backup, NetBackup uses a single exclude list—the list that contains the most specific name. For example, if there are files named: <pre> exclude_list.wkstations and exclude_list.wkstations.fulls </pre> NetBackup uses only: <pre> exclude_list.wkstations.fulls </pre>

Files that are excluded by Microsoft Windows Backup

Windows maintains a list of files and folders that are excluded when Microsoft Windows Backup is used to back up files. This list is known as the **FilesNotToBackup** list. NetBackup excludes those files and directories from automatic backups even if they are not in the NetBackup exclude list for the client. Those items also are excluded from user-directed backups (unlike items in a NetBackup exclude list, which can be backed up by a user-directed operation).

Windows also maintains a list of registry keys that are not to be restored. NetBackup does not restore the registry keys that are listed in the **Windows KeysNotToRestore** list.

Disaster Recovery tab

The **Disaster Recovery** tab appears for the **NBU-Catalog** policy type. The **Disaster Recovery** tab contains options for configuring disaster recovery protection methods for the catalog data.

Note: Do not save the disaster recovery information to the local computer. It is recommended to save the image file to a network share or a removable device.

Figure 20-31 Disaster Recovery tab




Table 20-53 describes the options on the **Disaster Recovery** tab.

Table 20-53 Options on the Disaster Recovery tab

Option	Description
Path	Browse and specify the directory where the disaster recovery information is to be saved. Do not save the disaster recovery information to the local computer. It is recommended that you save the image file to a network share or a removable device. The share must be established and available before the hot catalog backup runs. Specify an NFS share or a UNC path (CIFS Windows share). Note: The path cannot contain non-ASCII characters. When indicating a UNC path, note the following: ■ A Windows primary server can indicate a UNC path to a Windows computer. ■ A UNIX primary server cannot indicate a UNC path to a Windows computer. ■ A UNIX primary server cannot indicate a UNC path to a UNIX computer. To do so, first mount the UNC location on the primary server, and then provide the UNC path to the UNIX computer. On UNIX: The path for the disaster recovery information cannot be to a directory that is on the same partition as <code>/usr/openv/netbackup</code>. If the path is to a location on the same partition as <code>/usr/openv/netbackup</code>, NetBackup displays a status 20 error message. The message states that the disk path is invalid. Change the path on the Disaster Recovery tab to a directory on a different partition.
Logon	Specify the logon and password information that is required to access an established Windows or NFS share. If the logon information is not valid, NetBackup returns a message. The message requests that the user either reenter the logon and password information or clear the alternate location option to continue.
Password	Specify the password that is required to log on to the share.

Table 20-53 Options on the Disaster Recovery tab (*continued*)

Option	Description
Send in an email attachment	Specify the email address where the disaster recovery report should be sent. It is recommended that the disaster recovery report be sent to at least one email address. To send the information to more than one address, separate email addresses with a comma as follows: <i>email1@domain.com,email2@domain.com</i> On Windows: The <code>nbmail.cmd</code> or <code>mail_dr_info.cmd</code> script must be configured (<i>Install_path\NetBackup\bin\goodies\</i>). In addition specify the email addresses in the Disaster Recovery tab. On Windows: NetBackup performs the notification by passing the email addresses, subject, and message to <code>nbmail.cmd</code> or <code>mail_dr_info.cmd</code>. The scripts use the mail program that is specified in the script to send email to the user. See the comments in the script for configuration instructions. On Windows: The following points describe how <code>mail_dr_info.cmd</code> and <code>nbmail.cmd</code> interact: ■ If <i>Install_path\NetBackup\bin\mail_dr_info.cmd</i> is configured, the disaster recovery report is sent to the email address of the administrators that are indicated in the Disaster Recovery tab. NetBackup administrators can set up the script to send the disaster recovery information to alternate locations. ■ If <code>mail_dr_info.cmd</code> is not configured, and <i>Install_path\NetBackup\bin\goodies\nbmail.cmd</i> is configured, the disaster recovery report is sent to the administrators that are indicated in the Disaster Recovery tab by <code>nbmail.cmd</code>. ■ If neither file is configured, NetBackup attempts to use Microsoft internal IMAPI services. Note: On Windows: By default, neither <code>nbmail.cmd</code> nor <code>mail_dr_info.cmd</code> is configured to send email. See “Configure the nbmail.cmd script on the Windows hosts” on page 1048. On Windows: For more information on <code>mail_dr_info.cmd</code>, see the NetBackup Administrator's Guide, Volume II.
Critical policies	Lists the policies that are considered crucial to the recovery of a site in the event of a disaster. The NetBackup Disaster Recovery report lists all of the media that is used for backups of critical policies, including the most recent full backup. The NetBackup Disaster Recovery wizard warns you if any media for critical policies are not available. Note: The Disaster Recovery report lists the media for only incremental and full backup schedules so critical policies should use only incremental or full backup schedules. Certain database backups schedules, such as Oracle and Microsoft SQL Server, only use schedule types of Application Backup and Automatic Backup. Because of the schedule types, media listings for these backups do not appear on the Disaster Recovery report.


Note: Vault protects the disaster recovery data by sending the data to the Vault site as an email attachment of the Vault report email.

Adding policies to the Critical Policies list of a catalog backup policy

Use the following procedure to add policies to the **Critical Policies** list of a catalog backup policy.

To add a policy to the critical policies list

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Do one of the following:
 - Double-click a configured catalog backup policy.
 - Create a catalog backup policy.
- 3 Select the **Disaster Recovery** tab.
- 4 Near the **Critical Policies** list, click **Add**. An active field appears in the list.



- 5 Click the icon at the far right of the active field to display a list of configured policies. Select a policy to add to the **Critical Policies** list.
- 6 Do any of the following:

To add another policy	Click Add .
To change a policy	Select the policy and click Change .
To delete a policy	Select the policy and click Delete .
- 7 Click **OK** to save the **Critical policies** list and the other settings on the **Disaster Recovery** tab.

Creating a Vault policy

A Vault policy differs from other policies in the following respects:

- **Vault** must be specified as the policy type.
- No clients are specified in Vault policies, so the **Clients** tab does not appear.
- In the **Backup Selections** list, a Vault command is specified instead of files.

To create a Vault policy

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2** On the **Actions** menu, click **New > Policy**.
- 3** Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.
- 4** On the **Attributes** tab, select **Vault** as the policy type.
- 5** On the **Schedules** tab, click **New** to create a new schedule. The type of backup defaults to **Automatic**.

The **Clients** tab does not appear for Vault policy types.

- 6** Complete the schedule.
- 7** On the **Backup Selections** tab, enter one of two Vault commands:

`vltrun` Use `vltrun` to specify the robot, vault name, and profile for the job. The `vltrun` command accomplishes all the steps necessary to select, copy, and eject media. If the vault profile name is unique, use the following format:

```
vltrun
  profile_name
```

If the vault profile name is not unique, use the following format:

```
vltrun
  robot_number/vault_name/profile_name
```

`vlteject` Use the `vlteject` command to eject media or to generate reports for completed Vault sessions. For example:

```
vlteject -eject -report [-vault
vault_name

[-sessionid
id
]] [-auto y|n] [-eject_delay
seconds
]
```

Both commands are located in the following directory:

- On Windows:
`install_path\NetBackup\bin`
- On UNIX:
`/usr/opensv/netbackup/bin/`

For more information on Vault names, profile names, and command usage, see the [Vault Administrator's Guide](#).

8 Click **OK**.

Creating a BigData policy

Use the BigData policy to backup big data applications or certain hyper converged systems. For example, Hadoop or Nutanix Acropolis Hypervisor, respectively.

A BigData policy differs from other policies in the following respects:

- You must specify **BigData** as the policy type.
- The entries provided in the **Clients** tab and the **Backup Selections** tab differ based on the application that you want to back up.
- In the **Backup Selections** tab, you must specify certain parameters and their appropriate values.

To create a BigData policy

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2** On the **Actions** menu, click **New > Policy**.
- 3** Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.

- 4 On the **Attributes** tab, select **BigData** as the policy type.
- 5 On the **Schedules** tab, click **New** to create a new schedule. The type of backup defaults to **Automatic**.

Note: Currently, certain big data applications do not support all schedule types. For example, Nutanix supports only full backups.

- 6 Complete the schedule.
- 7 On the **Clients** tab, enter appropriate values according to your application type.
- 8 On the **Backup Selections** tab, enter appropriate parameters according to your application type.
- 9 Click **OK**.

NetBackup provides support to back up the following applications.

- Nutanix Acropolis Cluster
- Hadoop

To back up a Nutanix Acropolis Cluster

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.
- 4 On the **Attributes** tab, select **BigData** as the policy type.
- 5 On the **Schedules** tab, click **New** to create a new schedule. Currently, NetBackup supports only full backups for a Nutanix Acropolis Cluster.
- 6 On the **Clients** tab, enter the display name of the virtual machine.
- 7 On the **Backup Selections** tab, enter the following parameters and their values as shown:
 - `Application_Type=Nutanix-AHV`
The parameter values are case-sensitive.
 - `Backup_Host=<IP address or the hostname of the backup host>`
The backup host must be a Linux machine. The backup host can be a NetBackup client or a media server.

- `Application_Server=<IP address or the hostname of the Nutanix cluster>`

8 Click **OK** to save the changes.

To back up a Hadoop cluster

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2** On the **Actions** menu, click **New > Policy**.
- 3** Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.
- 4** On the **Attributes** tab, select **BigData** as the policy type.
- 5** On the **Schedules** tab, click **New** to create a new schedule.
- 6** On the **Clients** tab, enter the Fully Qualified Domain Name (FQDN) of the Hadoop cluster or name node.
- 7** On the **Backup Selections** tab, enter the following parameters and their values as shown:
 - `Application_Type=hadoop`
The parameter values are case-sensitive.
 - `Backup_Host=<IP_address or hostname>`
The backup host must be a Linux machine. The backup host can be a NetBackup client or a media server.
 - File path or the directory to back up.
- 8** Click **OK** to save the changes.

Performing manual backups

A manual backup is user-initiated and is based on a policy.

A manual backup is useful in the following situations:

- To test a configuration
- To back up a client that missed the regular backup
- To back up a client before new software is installed to preserve the old configuration
- To preserve records before a special event such as a company split or merger
- To back up quarterly or yearly financial information

In some cases, it may be useful to create a policy and schedule that is used only for manual backups. Create a policy for manual backups by creating a policy with a single schedule that has no backup window. Without a backup window, the policy can never run automatically.

To perform a manual backup

- 1 On the left select **Protection > Policies**.
- 2 Select the check box for the policy.
- 3 Select **Manual backup**. (To perform a manual backup, you must enable the **Go into effect at** attribute.)

See [“Go into effect at \(policy attribute\)”](#) on page 681.

If the **Go into effect at** attribute is set for a future date and time, the backup does not run.

- 4 Select the schedule and the clients that you want to back up.

If you do not select any schedules, NetBackup uses the schedule with the highest retention level. If you do not select any clients, NetBackup backs up all clients.

User schedules do not appear in the schedules list. A user schedule cannot be manually backed up because it does not have a backup selection list (the user selects the files).

- 5 Select **Backup** to start the backup.

Active Directory granular backups and recovery

Administrators can use NetBackup to restore individual objects and attributes in the Active Directory instead of restoring the entire Active Directory. Administrators can also restore deleted objects (tombstone objects) from the Active Directory.

The following topics describe how to configure a policy to perform recovery of an Active Directory object:

- System requirements necessary to perform Active Directory granular backups and restores.
- How to configure a policy for an Active Directory backup that allows granular restores.
- How to restore individual objects and attributes in the Active Directory.

System requirements for Active Directory granular NetBackup backups and recovery

For a list of operating system versions and media server platforms that support Active Directory granular restore, see the *NetBackup Enterprise Server and Server - Software Compatibility List* at the following URL:

https://www.veritas.com/support/en_US/article.100040093

To perform Active Directory granular backups and restores, ensure that you meet the following requirements:

- The Network File System (NFS) must be installed on the media server and all Active Directory domain controllers or ADAM/LDS hosts.
See [“About installing and configuring Network File System \(NFS\) for Active Directory Granular Recovery”](#) on page 1133.
See [“About configuring Services for Network File System \(NFS\)”](#) on page 1134.
- The NetBackup Client Service must be configured to log on as an account with domain privileges.
To perform granular backups and restores of the Active Directory, the NetBackup Legacy Client Service (`bpinetd`) must run under the domain administrator account on the Active Directory domain controller or ADAM server. By default, `bpinetd` runs under the Local System account.
See [“Configuring the NetBackup Client Service”](#) on page 1052.

Creating a policy that allows Active Directory granular restores

A NetBackup policy that backs up the Active Directory can be configured to allow the restore of the objects and attributes in the Active Directory. The objects and attributes can be restored locally or remotely without the interruption of restarting the domain controllers where the restore is performed.

The **Active Directory** host properties determine if NetBackup performs a consistency check when Microsoft Volume Shadow Copy Service (VSS) is used as the snapshot provider.

See [“Active Directory properties”](#) on page 55.

To create a policy to allow Active Directory restores

- 1 Check that the NetBackup Legacy Client Service (`bpineta`) is running under the domain administrator account on the Active Directory domain controller. In this case, the Active Directory domain controller is the NetBackup client.

See [“Configuring the NetBackup Client Service”](#) on page 1052.
- 2 In the **Policy** dialog box, on the **Attributes** tab, select **MS-Windows** as the policy type. Specify the other policy attributes as needed.
- 3 Enable the **Enable granular recovery** option. If this option is not enabled, the backup still runs, but the backup cannot produce granular restores.
- 4 In the **Schedules** tab, create schedules as needed.

Other items in the policy may use a differential or cumulative incremental backup type, but the Active Directory items are always fully backed up.

See [“Active Directory backups are full backups”](#) on page 837.
- 5 In the **Backup Selections** tab, open the **Select Directive** dialog.
- 6 For the **Directive set**, select **Windows 2003** or **Windows 2008**.
- 7 To back up the Active Directory, select any one of the following directives:
 - See [“System State:\ directive”](#) on page 816.
 - See [“Shadow Copy Components:\ directive”](#) on page 816.
 - See [“ALL_LOCAL_DRIVES directive”](#) on page 814.

Note: Active Directory Application Mode (ADAM) is a lightweight directory service that runs as a user service. This directive can be used to back up ADAM data on computers where it is installed. However, it does not back up the Active Directory itself.

- 8 In the **Clients** tab, select the clients as needed.
- 9 Save the policy.

Active Directory backups are full backups

Any Active Directory backup is always a NetBackup full backup, whether it is a granular backup or not.

Whenever Active Directory is in a policy’s **Backup Selections** list, the Active Directory portion is always fully backed up, even when the backup type is incremental, differential or cumulative. Any other items in the **Backup Selections** list may use a differential or cumulative incremental backup type as indicated. Even


though a full backup is forced for an Active Directory backup, normal incremental rules are applied to the non-Active Directory items in the policy file list.

Restoring Active Directory objects

The following procedure describes how to restore objects from an Active Directory backup in a non-disaster recovery situation:

To restore individual objects from an Active Directory backup

- 1 Open the NetBackup Backup, Archive, and Restore client interface.
- 2 Select **File > Select Files and Folders to Restore**.
- 3 Expand and browse the **Active Directory** node.
- 4 Select the objects to be restored. Do not select both granular and non-granular objects. When a user explores and expands selections, a delay can occur during communication with the NetBackup server. The delay is a result of dynamically determining the contents from the image on the media server. The approach prevents the NetBackup catalog from unanticipated growth due to numerous granular entries.



- 5 Select **Action > Restore**.

- 6 If an Active Directory object is selected, the **Restore Marked Files** dialog box contains two tabs:
 - **General tab**

When an Active Directory object is selected, the **Restore Destination Choices** are disabled in the **General** tab. Configure the other restore options as needed.
 - **Active Directory tab**

The **Active Directory** tab contains an option to recreate the objects that have been deleted: **Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container**.

The **Active Directory** tab contains an option that lets administrators recreate the objects whose tombstone lifetimes have passed. The objects have also been purged from the Active Directory Deleted Objects container.

To allow this capability, enable the option labeled **Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container**.
- 7 Click **Start Restore** in the **Restore Marked Files** dialog box.

Some restore situations require additional steps, depending on what is restored.

See [“Troubleshooting granular restore issues”](#) on page 839.

Troubleshooting granular restore issues

Some granular restore situations require additional steps to fully restore the objects. In other situations, a granular restore of some part of the Active Directory is not possible.

[Table 20-54](#) describes potential problems for granular restores.

Table 20-54 Troubleshooting restore issues

Situation	Recommendation
Restores that are disabled	When user and computer accounts are restored from a granular Active Directory restore, they are sometimes disabled. The following are possible reasons why the accounts can be disabled: ■ When objects in Active Directory are deleted, they are removed from their current Active Directory or ADAM/AD LDS container. They are converted into tombstones and placed in the Active Directory Deleted Objects container where their tombstone lifetime is monitored. By default, NetBackup restores deleted objects from this container if the tombstone lifetime has not passed. After the tombstone lifetime passes, the tombstones are purged from the Active Directory Deleted Objects container. Purging the tombstones has the effect of permanently deleting the objects from the Active Directory and ADAM/AD LDS databases. ■ When restoring user objects, you must reset the object's user password and enable the object's user account: ■ For Active Directory user objects, use the Microsoft Active Directory Users and Computers application. ■ For ADAM/AD LDS user objects, use ADSI Edit. In Active Directory, computer objects are derived from user objects. Some attributes that are associated with a computer object cannot be restored when you restore a deleted computer object. They can only be restored if the attributes were saved through schema changes when the computer object was originally deleted. ■ Computer object credentials change every 30 days and the credentials from the backup may not match the credentials that are stored on the actual computer. When a computer object is restored it is disabled if the userAccountControl property was not preserved in the deleted object. Use the Microsoft Active Directory Users and Computers application to reset the account of a computer object: ■ Remove the computer from the domain. ■ Re-join the computer to the domain. The security identifiers (SID) for the computer remains the same since it is preserved when a computer object is deleted. However, if the tombstone expired and a new computer object was recreated, the SID is different.

Table 20-54 Troubleshooting restore issues *(continued)*

Situation	Recommendation
Group and member objects	To restore Active Directory group membership links may require that the restore job be run twice. For example, consider the case where a group and its member objects are deleted. If a restore job contains both group objects and member objects, the job restores the objects in alphabetical order. However, the group that is restored has a link dependency on a member that does not exist yet. When the group is restored, the link cannot be restored. Run the restore again to restore all forward and backward links.
Group policy objects	NetBackup does not support granular restores of Group Policy Objects.

Synthetic backups

This chapter includes the following topics:

- [About synthetic backups](#)
- [Recommendations for synthetic backups and restores](#)
- [Synthetic full backups](#)
- [Synthetic cumulative incremental backups](#)
- [Schedules that must appear in a policy for synthetic backups](#)
- [Adding clients to a policy for synthetic backups](#)
- [Change journal and synthesized backups](#)
- [True image restore and synthesized backups](#)
- [Displaying synthetic backups in the Activity Monitor](#)
- [Logs produced during synthetic backups](#)
- [Synthetic backups and directory and file attributes](#)
- [Using the multiple copy synthetic backups method](#)
- [Optimized synthetic backups](#)
- [Optimized synthetic backups for deduplication](#)

About synthetic backups

During a traditional full backup, all files are copied from the client to a primary server or a media server. The files are copied even though those files may not have changed since the last incremental backup.

When NetBackup creates a synthetic full backup, NetBackup detects whether new or changed files have been copied to the media server during the last incremental backup. The client does not need to be running to combine the full backups and the incremental backups on the media server to form a new, full backup. The new, full synthetic backup is an accurate representation of the clients' file systems at the time of the most recent full backup.

Because processing takes place on primary and media servers instead of the client, synthetic backups help to reduce the network traffic. Files are transferred over the network only once. After the backup images are combined into a synthetic backup, the tapes or disk that contain the component images can be recycled or reclaimed. Synthetic backups can reduce the number of tapes or disk space in use.

Synthetic backups can be written to tape storage units or disk storage units, or a combination of both. If the backups use tape, the backups can be synthesized when drives are not generally in use. For example, if backups occur primarily at night, the drives can synthesize full backups during the day.

The **Synthetic Backup** option is available under the following conditions:

- The policy type must be either Standard or MS-Windows.
- The **Collect True Image Restore Information With Move Detection** option must be selected on the **Policy Attributes** tab.
 See ["Collect true image restore information \(policy attribute\) with and without move detection"](#) on page 693.
- The schedule that is created for a synthetic backup must have **Synthetic Backup** selected.
 See ["Synthetic backup \(schedule attribute\)"](#) on page 741.
- One of the following must be available:
 - Disk storage unit(s) with adequate space available.
 - Tape library(s) with multiple drives to read and write.
 See ["Recommendations for synthetic backups and restores"](#) on page 843.
 - A combination of disk storage unit(s) and tape library(s).

Recommendations for synthetic backups and restores

The synthetic full backup is a scalable solution for backing up remote offices with manageable data volumes and low levels of daily change.

If the clients experience a high rate of change daily, the incremental backups are too large. In this case, a synthetic backup is no more helpful than a traditional full backup.

Synthetic backups are supported on all media server platforms and tier one primary server platforms.

The following items describe recommendations to use synthesized backups to full advantage, and situations under which synthesized backups are not supported:

Recommendations concerning backups:

- Do not multiplex any backups that are to be synthesized because it is inefficient. To synthesize multiplexed client images requires multiple passes over the source media—one per client.
 Performance issues can also occur if multiple streams are selected for synthesized backups. The issues are similar to those encountered while multiplexing synthesized backups. Back up to disk whenever possible to improve multiple stream performance issues.
- It is recommended that you not enable **Expire after copy** retention for any storage units that are to be used with SLPs with either of the following: Accelerator or synthetic backups. The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.
- Reduce the gap between the last incremental backup and the synthesized backup. Since a synthetic backup does not involve direct contact with the client, a synthetic backup is only as current as the last incremental backup. If there is a concern to reduce a potential gap in backup coverage, run an incremental backup before the synthetic backup.
- You can create multiple copies with synthetic backups by using the multiple copies synthetic backup method.
 Although synthetic backups do support the use of storage lifecycle policies, SLPs cannot be used for the multiple copy synthetic backups method.
 See [“Using the multiple copy synthetic backups method”](#) on page 853.
- Synthetic backups are not supported if any of the component images are encrypted.
- A user-generated backup cannot be used to generate a synthetic image. A backup that is generated from a User Backup schedule or a User Archive schedule cannot be used as one of the components of a synthetic backup.
- Synthetic backups and optimized synthetic backups do not support Auto Image Replication.

Recommendations concerning restores:

- The time that is required to perform a restore from a synthetic backup does not increase significantly over time.
- The restore times for both a complete synthetic backup and for a single file is the same. It is the same whether the restore is from a traditional backup or from a synthetic backup.
- The restore time of a single directory may increase over time when sourced from synthetic backups. The restore time depends on the pattern of file changes within the directory.
- Contrast a traditional full backup, which stores the files in file system order with a synthetic full backup, which stores the files in last-file-accessed order. The synthetic full contains the newest files at the front of the media and the unchanged files at the end. Over time, the processing order introduces the potential for fragmentation of a single directory across the synthetic full image.
- Note that the scenario is limited to single directory restores. Single file restores and full image restores from synthetic fulls are equal or better than from traditional full backups, as noted in previous bullets.
- If checkpoint restart is indicated for the policy, the backups that are produced with the synthetic backup schedule are not checkpointed. The option is enabled if **Take checkpoints** on the policy Attributes tab is enabled. If the **Take checkpoints** option is enabled for a synthetic backup, the property has no effect.

Table 21-1 Recommendations when using disk storage or tape storage for synthetic backups

Storage unit type	Recommendations
Disk storage units	Disk-based images are more efficient for synthesizing. NetBackup processes the newest component images first in a synthesized backup, followed by sequentially older images. When two or more component images are written to the same tape, the tape movement can be inefficient compared to disk-based images. Synthetic full backups are generated more quickly when built from disk-based incremental backups. If the synthetic full backup is also generated on disk, the run time is even faster. The disk copy then can be duplicated to tape.

Table 21-1 Recommendations when using disk storage or tape storage for synthetic backups *(continued)*

Storage unit type	Recommendations
Tape storage units	If tape is used instead of disk, the tape for the synthetic image must be different from the tape where the component images reside. The maximum drive usage applies only to the drive that is needed for writing the synthetic backup. If any of the component images reside on tape, an additional drive is needed for reading. If a single tape drive device is used to generate synthetic images, place component images in a hard drive location first. In that way, a synthetic image can be generated with the single tape drive device.


Synthetic full backups

A synthetic backup can be a synthetic full or a synthetic cumulative backup.

The images that are used to create the synthetic image are known as component images. For instance, the component images in a synthetic full are the previous full image and the subsequent incremental images.

[Figure 21-1](#) illustrates the creation of synthetic full backups (B, C, D) from an existing full backup (A) and shows the incremental backups between full backups.

Figure 21-1 Creation of synthetic full backups



The traditional full backup (A) and the incremental backups are created in the traditional manner: data is scanned, and then copied from the client's file system to the backup media. The synthetic backups do not interact with the client system at all, but are instead synthesized on the media server.

See ["Synthetic cumulative incremental backups"](#) on page 848.

The following is an example of a synthetic full backup:

- Create a Standard or MS-Windows policy for the clients you want to back up. Include the following schedules:
 - A schedule for one full, traditional backup to run at least once.
 - A schedule for daily (Monday through Saturday) differential incremental backups.
 - A schedule for weekly full, synthetic backups.
- Make sure that the traditional full backup runs. If the backup does not complete, run the backup manually.

- Per schedule, run daily, differential incremental backups for the clients throughout the week. The last incremental backup for the week runs on Saturday.
- Per schedule, run synthetic full backups for the clients on subsequent Sundays.

Note: The synthetic full backups in the scenario are only as current as the Saturday incremental backup.

Synthetic cumulative incremental backups


The scenario to create a synthetic, cumulative incremental backup is similar to the scenario to create a synthetic full backup. Remember, a cumulative incremental backup includes all changes since the last full backup.

If a cumulative incremental backup exists that is newer than the last full backup, a synthetic cumulative backup image is produced by consolidating the following component backup images:

- All differential incremental backups that were taken since the last cumulative backup.
- The last cumulative incremental backup. If no cumulative incremental backup is available, only the differential incremental backups are used for the synthetic image.

[Figure 21-2](#) illustrates the creation of synthetic cumulative incremental backups (A, B, C) from the latest cumulative incremental backup and shows the subsequent differential incremental backups.

Figure 21-2 Creation of synthetic cumulative backups



The following is an example of a synthetic cumulative backup:

- Create a Standard or MS-Windows policy for the clients (5.0 or later) you want to back up. Include the following schedules:
 - A schedule for one full, traditional backup to run at least once.
 - A schedule for daily (Monday through Saturday) differential incremental backups.
 - A schedule for weekly cumulative incremental synthetic backups.
- Make certain that the traditional full backup runs. If the backup does not complete, run the backup manually.
- Per schedule, run daily differential incremental backups for the clients throughout the week. The last incremental for the week runs on Saturday.
- Per schedule, run synthetic cumulative incremental backups for the clients on subsequent Sundays.

Note: The synthetic cumulative backups in the scenario are only as current as the Saturday incremental backup.

Schedules that must appear in a policy for synthetic backups

A policy for synthetic backups must contain one of the following types of schedules:

- At least one traditional, full backup must be run successfully to create a full image. The synthetic backup job fails if there is not at least one previous full image.
- Schedule(s) for incremental backups.
 Incremental backups are necessary to capture the changes in the file system since the last full or incremental backup. The synthetic backup job receives a status code of 1 for a policy that contains full or incremental synthetic backup schedules, but no incremental backup schedules.
 The synthetic backup synthesizes all of the incremental backups to create a new full or cumulative backup image. Therefore, the synthetic backup is only as current as the last incremental backup.

Note: To configure a synthetic cumulative backup for any clients that are archive bit-based (default), use only differential incremental backups for the traditional, non-synthesized backups.

- One full and one cumulative backup schedule with the **Synthetic Backup** option selected.
 See [“Synthetic backup \(schedule attribute\)”](#) on page 741.

Adding clients to a policy for synthetic backups

After clients are added to a synthetic backup policy, run a traditional, full backup of the policy. A traditional backup is necessary before a synthetic backup can be created.

Since **Collect True Image Restore Information With Move Detection** is required for synthetic backups, all of the clients in the policy must support TIR.

See [“Collect true image restore information \(policy attribute\) with and without move detection”](#) on page 693.

Change journal and synthesized backups

If the **Use Change Journal** host property on a Windows client is enabled, the property has no effect when the client is backed up using the synthetic backup schedule.

See [“Client settings properties for Windows clients”](#) on page 77.

True image restore and synthesized backups

Since the **Collect true Image restore information with move detection** policy property must be enabled for synthetic backups, all clients that are included in the policy must support TIR.

See [“Collect true image restore information \(policy attribute\) with and without move detection”](#) on page 693.

The **Keep true image restoration (TIR) information** property indicates how long TIR information in the image catalog is kept before it is pruned (removed). The property is located in the primary server **Clean-Up** host properties.

See [“Clean up properties”](#) on page 61.

However, if a synthetic full and synthetic cumulative schedule was defined in the policy, the TIR information is pruned from the component images until a subsequent traditional or synthetic full or cumulative backup image has generated successfully.

Consider a situation where **Keep true image restoration (TIR) information** host specifies that TIR information is pruned from the catalog after two days. On the third day the TIR information is pruned only if a traditional or synthetic full backup image has been generated.

If the TIR information was pruned from a component image and you accidentally expire the most recent synthetic image, rerun the synthetic backup job to restore automatically the TIR information to the catalog. In case the TIR information cannot be restored due to bad, missing, or vaulted media, the synthetic backup job fails with error code 136 (TIR info was pruned from the image file). If the problem is correctable, run the synthetic backup again.

Displaying synthetic backups in the Activity Monitor

A synthetic job is distinguished from a traditional full backup by the notation that is indicated in the Data Movement field of the Activity Monitor. Synthetic jobs display Synthetic as the Data Movement type while traditional backups display Standard.

Logs produced during synthetic backups

When a synthetic backup is scheduled, NetBackup starts the `bpsynth` program to manage the synthetic backup process. `bpsynth` plans how the synthetic backup is built from the previous backup images.

If it is needed, `bpsynth` schedules the tape drive resources that are needed for the synthetic backup. If the required resources are not available, the job fails with a status code that indicates that a resource is needed.

If the resources can be obtained eventually but not immediately, the synthetic job waits until the resources become available. A synthetic job may wait while a backup, restore, or another synthetic backup job uses a drive.

`bpsynth` passes the information to programs `bptm` and `bpdm` so that tape and disk images can be read or written. Catalog information is managed using `bpdbm`. Each of these programs has a debug log file in the logs directory.

If problems occur with synthetic backups, the following debug logs are required to diagnose the problem:

- On the primary server: `bpsynth`, `bpdbm`, and `vnetd`.
- On the media server(s): `bpcd`, `bptm` (if any images are written to or from a tape device), and `bpdm` (if any images are written to or from disk).
Note that several media servers can be involved if the component images are on different nodes.

The log files for synthetic backups are described in the [NetBackup Logging Reference Guide](#).

However, `bpsynth` is used for each stream or client. To use `bpsynth` can be inefficient with tape images since `bpsynth` needs a tape drive to write the new image. Also, `bpsynth` may use the same component image volumes. One may need to finish before the next can proceed.

Synthetic backups and directory and file attributes

For a synthetic backup to include directory and the file attribute changes, the change must first be picked up by a component incremental backup. (For example, changes like Access Control Lists (ACLs).)

On UNIX, changing an object's ACL changes the `ctime` (inode change time) for the object but not the `mtime` (data modification time). Since `mtime` starts incremental backups, the ACL change is not reflected in an incremental backup, and therefore not in a synthetic full backup.

To include ACL changes in backups, enter `USE_CTIME_FOR_INCREMENTALS` in the `bp.conf` file on each UNIX client.

For each Windows client, enable **Based on archive bit** (in the **Incrementals** settings). The property is found in the **Windows client > Client settings** host properties.

See [“Client settings properties for Windows clients”](#) on page 77.

Using the multiple copy synthetic backups method

The multiple copy synthetic backups method introduces the capability to produce a second copy of a synthetic backup at a remote site as part of a normal synthetic backup job.

This method provides the following benefits:

- It eliminates the bandwidth cost of copying synthetic full backups to another site.
Instead of duplicating a local synthetic full backup to a remote site to produce a second copy, it is more efficient to produce the second copy by using data movements only at the remote site.
- It provides an efficient method to establish a dual-copy disaster recovery scheme for NetBackup backup images.

[Table 21-2](#) emphasizes how the synthetic full backup produced at the remote site is a clone, or a second copy, of the first copy produced at the local site.

Table 21-2 Comparing synthetic copy process with and without method enabled


Step	Without using the multiple copy synthetic backups method:	Using the multiple copy synthetic backups method:
1	A full backup is performed at the local site (Site A).	Step 1 remains the same.
2	The full backup is duplicated to the remote site (Site B).	Step 2 remains the same.
3	An incremental backup is performed at Site A.	Step 3 remains the same.
4	The incremental backup is duplicated to Site B.	Step 4 remains the same.
5	Steps 3 and 4 are repeated each time an incremental schedule runs.	Step 5 remains the same.
6	A full synthetic backup is produced at Site A.	Step 6 remains the same.

Table 21-2 Comparing synthetic copy process with and without method enabled (*continued*)

Step	Without using the multiple copy synthetic backups method:	Using the multiple copy synthetic backups method:
7	The full backup is duplicated to Site B.	A full synthetic backup is produced at Site B from images at Site B. The full synthetic backup at the remote site is a second copy of the synthetic backup at the local site.
8	Steps 2 through 7 repeat per backup scheduling needs.	Step 8 remains the same.

Figure 21-3 shows how no extra bandwidth is used to copy the synthetic full backup from Site A to Site B.

Figure 21-3 Remote creation of synthetic full backup



Configuring multiple copy synthetic backups

To configure a multiple copy synthetic backup, create a configuration file on the primary server for each synthetic backup policy for which a second copy is to be produced.

The configuration file is a text file that is named after the policy and schedule:

`multi_synth.policy.schedule`

Create the file in the following location:

- On Windows:
`install_path\NetBackup\db\config\multi_synth.policy.schedule`
- On UNIX:

`/usr/opensv/netbackup/db/config/multi_synth.policy.schedule`

Configuration variables for multiple copy synthetic backups

The file format uses a traditional name-pair scheme for setting configuration preferences. Each preference uses a key name that is separated from the preference value by an equal sign with each name-value pair residing on a single line.

For example:

```
NAME=VALUE
```

Enter all values as integers.

[Table 21-3](#) describes the configuration entries that can be included in the configuration file.

Table 21-3 Configuration entries

Entry	Purpose
SRC_COPY	Specifies the copy number of each source component for the second synthetic backup. Every source backup must have a copy by this number unless SRC_COPY_FALLBACK is specified. The default is 2.
TARGET_COPY	Specifies the copy number for the second synthetic backup produced. The TARGET_COPY value must be different from the copy number of the first synthetic backup (which is 1). Default is 2.
COPY	COPY is an alternate specification for SRC_COPY and TARGET_COPY. If COPY is specified and either SRC_COPY and TARGET_COPY is not specified, the value for COPY is used.
TARGET_STU	Specifies the storage unit name or storage unit group name where the second copy synthetic backup is to be written. Use the special identifier __ANY__ to indicate that Any Available storage unit can be used that is not configured to be on demand only. Note that there are two underscores before and after ANY: TARGET_STU=__ANY__
FAIL_MODE	The second synthetic backup is produced immediately following the first copy synthetic backup if no errors occur during production of the first copy. If an error occurs during the second copy, the FAIL_MODE value specifies the fate of the first copy job and image. Specify one of the following: - FAIL_MODE=ALL ALL means that if the second copy fails, the first copy and its job also fail. (Default.) - FAIL_MODE=ONE ONE means that if the second copy fails, the failure does not affect the first copy job.

Table 21-3 Configuration entries (*continued*)

Entry	Purpose
ENABLED	Specifies whether production of the second copy is enabled or disabled. This entry turns on the feature. Specify one of the following: - ENABLED=YES Production of the second copy is enabled. (Default.) - ENABLED=NO Production of the second copy is disabled.
SRC_COPY_FALLBACK	Specifies that if a copy by the number given in SRC_COPY or COPY does not exist, the synthetic backup should use the primary backup. The only valid value is the following: SRC_COPY_FALLBACK=PRIMARY
VOLUME_POOL	Specifies the volume pool for tape media, if one is used. If no volume pool is specified, NetBackup uses the volume pool that is specified in the policy. If a volume pool is entered for disk, the entry is ignored.

Multiple copy synthetic backups configuration examples

The following multiple copy synthetic configuration example takes advantage of default values to produce the second synthetic copy.

```
TARGET_STU=disk_stu
```

The default source of copy 2 and the default destination copy 2.

In this example, the second copy targets a tape library (`tape_stu`). The configuration specifies a volume pool (`Synthetics`) for the target copy.

The copy number for the multiple copy synthetic backup is copy 3. If copy 3 is unavailable, `SOURCE_COPY_FALLBACK` indicates that copy 3 can be produced using the primary copy.

If copy 3 fails, only copy 3 fails and not the job of the primary copy.

```
TARGET_STU=tape_stu
VOLUME_POOL=Synthetics
SOURCE_COPY_FALLBACK=PRIMARY
COPY=3
ENABLED=YES
FAIL_MODE=ONE
```

Optimized synthetic backups

NetBackup environments that use the Data Protection Optimization Option license can benefit from the optimized synthetic backup method. Optimized synthetic backups take advantage of the capabilities of the OpenStorage API.

This method constructs the synthetic image by using calls from the media server to the storage server. The media server tells the storage server which full and incremental images to use to create the synthetic backup. Then, the storage server constructs (or synthesizes) the synthetic image directly on the storage server, reducing network traffic.

For more information, see the following guides:

- [NetBackup Deduplication Guide](#)
- [NetBackup OpenStorage Solutions Guide for Disk](#)

Optimized synthetic backups for deduplication

NetBackup environments that use the NetBackup Data Protection Optimization Option license can benefit from the optimized synthetic backup method.

This method constructs the synthetic image by using calls from the backup server to the storage server. The backup server tells the storage server which full and incremental images to use to create the synthetic backup. Then, the storage server constructs (or synthesizes) the synthetic image directly on the storage server, reducing network traffic.

For more information, see the following guides:

- [NetBackup Deduplication Guide](#)
- [NetBackup OpenStorage Solutions Guide for Disk](#)

Protecting the NetBackup catalog

This chapter includes the following topics:

- [About the NetBackup catalog](#)
- [Parts of the NetBackup catalog](#)
- [Catalog backups](#)
- [Recovering the catalog](#)
- [Disaster recovery emails and the disaster recovery files](#)
- [Disaster recovery packages](#)
- [About disaster recovery settings](#)
- [Setting a passphrase to encrypt disaster recovery packages](#)
- [Archiving the catalog and restoring from the catalog archive](#)
- [Estimating catalog space requirements](#)

About the NetBackup catalog

A NetBackup catalog is the internal database that contains information about NetBackup backups and configuration. Backup information includes records of the files that have been backed up and the media on which the files are stored. The catalogs also contain information about the media and the storage devices.

Configure a disaster recovery pass phrase and a catalog backup before you run any regular backups. NetBackup needs information from the catalog to determine

where the backups of files are located. Without a catalog, NetBackup cannot restore data.

See [“Setting a passphrase to encrypt disaster recovery packages”](#) on page 875.

See [“Configuring catalog backups”](#) on page 867.

As additional protection for the catalog, consider archiving the catalog.

Parts of the NetBackup catalog

The NetBackup catalog resides on the NetBackup primary server. It manages and controls access to the following types of data:

- Image metadata (information about backup images and copies).
- Backup content data (information about the folders, files, and the objects in a backup (.*z* files)).
- NetBackup backup policies.
- NetBackup licensing data.
- The NetBackup error log.
- The client database.
- Cloud configuration files.

See [“About the catalog backup of cloud configuration files”](#) on page 864.

The catalog consists of the following parts:

- NetBackup stores information in the NetBackup database (NBDB). The metadata includes information about the data that has been backed up, and about where the data is stored.
See [“NetBackup databases and configuration files”](#) on page 860.
- The image database.
The image database contains information about the data that has been backed up.
See [“About the NetBackup image database”](#) on page 862.
- NetBackup configuration files.
- The key management service (KMS) configuration files
For more details on the KMS configuration, see the [NetBackup Security and Encryption Guide](#).

NetBackup is sensitive to the location of the primary server components. Running any part of NetBackup (the binaries, the logs, the database, the images) on a

network share (NFS, for example) can affect performance of even normal operations. NetBackup can be CIFS-mounted on SAN or NAS storage as long as the average I/O service times remain less than 20 milliseconds.

The storage must also meet certain conditions to ensure data integrity in the NetBackup catalog.

- The order of file writes must be guaranteed.
- When a write request is issued, the write must complete to the physical storage. The write request must not merely be buffered when the SAN or the NAS returns from the write call.

See the following article for more information:

https://www.veritas.com/content/support/en_US/article.100023390

NetBackup databases and configuration files

The NetBackup catalog backup includes the NetBackup databases and the configuration files, as follows.

Databases

The NetBackup databases include the NBDB database and the NetBackup Authorization database (NBAZDB). If Bare Metal Restore is installed (optionally-licensed) there is also the BMRDB database.

The databases are located in the following directories:

`install_path\NetBackupDB\data`

`/usr/opensv/db/data/`

These directories contain the following subdirectories:

`\bmrdb\` or `/bmrdb/` (if BMR is installed)

`\nbazdb\` or `/nbazdb/` (NetBackup authorization)

`\nbdb\` or `/nbdb/` (contains both the NBDB and the EMM databases)

Configuration files

Warning: Do not edit the configuration files. NetBackup may not start if you change these files.

Note: The catalog backup process copies this data to `/usr/opensv/db/staging` and backs up the copy.

The following configuration files are created:

```
pgbouncer.ini  
pg_hba.conf  
pg_ident.conf  
postgresql.auto.conf  
postgresql.conf  
userlist.txt  
vxdbms.conf  
web.conf
```

Most of the configuration files are located in the following directories:

```
install_path\NetBackupDB\data\instance  
  
/usr/opensv/db/data/instance
```

`web.conf` is created in the following directories:

```
/usr/opensv/var/global/wsl/config  
  
install_path\NetBackup\var\global\wsl\config
```

About the Enterprise Media Manager (EMM)

The Enterprise Media Manager (EMM) is a NetBackup service that manages the device and the media information for NetBackup. The Enterprise Media Manager stores its managed information in a database that resides on the primary server. The NetBackup Resource Broker queries EMM to allocate storage units, drives (including drive paths), and media.

EMM contains the following information:

- Device attributes
- Robotic library and standalone drive residence attributes
- NDMP attributes
- Barcode rule attributes
- Volume pool attributes
- Tape attributes
- Media attributes

- Storage unit attributes
- Storage unit group attributes
- Hosts with assigned tape drives
- Media and device errors
- Disk pool and disk volume attributes
- Storage server attributes
- Log on credentials for storage servers, disk arrays, and NDMP hosts
- Fibre Transport attributes

EMM ensures consistency between drives, robotic libraries, storage units, media, and volume pools across multiple servers. EMM contains information for all media servers that share devices in a multiple server configuration. The NetBackup scheduling components use EMM information to select the server, drive path, and media for jobs.

About the NetBackup image database

The image database contains subdirectories for each client that is backed up by NetBackup, including the primary server and any media servers.

The image database is located in the following location:

- **Windows:** `\Program Files\Cohesity NetBackup\NetBackup\db\images`
- **UNIX:** `/usr/opensv/netbackup/db/images`

The image database contains the following files:

Image files	Files that store only backup set summary information.
<code>.lck</code> files	Used to prevent simultaneous updates on images.
Image <code>.f</code> files	Used to store the detailed information about each file backup.
<code>db_marker.txt</code>	Used to ensure that access to the <code>db</code> directory is valid when the NetBackup Database Manager starts up. Do not delete this file.

The image database is the largest part of the NetBackup catalog. It consumes about 99% of the total space that is required for the NetBackup catalog. While most of the subdirectories are relatively small in the NetBackup catalogs, `\images` (Windows) or `/images` (UNIX) can grow to hundreds of gigabytes. The image database on the primary server can grow too large to fit on a single tape. Image database growth depends on the number of clients, policy schedules, and the amount of data that is backed up.

See [“Estimating catalog space requirements”](#) on page 888.

If the image catalog becomes too large for the current location, consider moving it to a file system or disk partition that contains more space.

See [“Moving the image catalog”](#) on page 890.

The catalog conversion utility (`cat_convert`) can be used to convert `.f` files into a human-readable format.

Information about the `cat_convert` command is available in the *NetBackup Commands Reference Guide*.

About NetBackup image `.f` files

The binary catalog contains one or more image `.f` files. This type of file is also referred to as a “files” file. The image `.f` file may be large because it contains the detailed backup selection list for each file backup. Generally, image files range in size from 1 kilobyte to 10 gigabytes.

Note: You can use intelligent catalog archiving (ICA) to reduce the number of catalog `.f` files based on a specified retention period or file size.

See [“Enabling intelligent catalog archiving \(ICA\) to reduce the number of `.f` files”](#) on page 880.

ICA applies only to servers running NetBackup 11.1 and later using MSDP or MSDP Cloud storage.

The `.f` files are found in the following location:

Windows: `install_path\NetBackup\db\images\clientname\ctime`

UNIX: `/usr/opensv/netbackup/db/images/clientname/ctime/`

The file layout determines whether the catalog contains one `.f` file or many `.f` files. NetBackup configures the file layout automatically, based on the size of the binary catalog. NetBackup uses one of two layouts: single file layout or multiple file layout.

- Image `.f` file single file layout

NetBackup stores file information in a single image `.f` file if the information for the catalog is less than 100 megabytes.

When the backup file of one catalog backup is less than 100 megabytes, NetBackup stores the information in a single image `.f` file. The image `.f` file is always greater than or equal to 72 bytes, but less than 100 megabytes.

The following is a UNIX example of an `.f` file in a single file layout:

```
-rw----- 1 root other 979483 Aug 29 12:23 test_1030638194_FULL.f
```

- Image .f file multiple file layout

When the file information for one catalog backup is greater than 100 megabytes, the information is stored in multiple .f files: one main image .f file plus nine additional .f files.

Separating the additional .f files from the image .f file and storing the files in the `catstore` directory improves performance while writing to the catalog.

The main image.f file is always exactly 72 bytes.

```
-rw- 1 root other      72 Aug 30 00:40 test_1030680524_INCR.f
-rw- 1 root other     804 Aug 30 00:08 catstore/test_1030680524_INCR.f-list
-rw- 1 root other 1489728 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgDir0
-rw- 1 root other      0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgExtraObj0
-rw- 1 root other 1280176 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgFile0
-rw- 1 root other     192 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgHeader0
-rw- 1 root other      0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgNDMP0
-rw- 1 root other 9112680 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgRecord0
-rw- 1 root other 2111864 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgStrings0
-rw- 1 root other      11 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgUserGroupNames0
```

About the catalog backup of cloud configuration files

The following cloud configuration files are backed up during the NetBackup catalog backup process:

All .txt files in the `meter` directory, which contain intermediate metering data

- CloudInstance.xml
- cloudstore.conf
- libstspienencrypt.conf
- libstspimetering.conf
- libstspithrottling.conf
- libstspicloud_provider_name.conf

All .conf files that are specific to the cloud providers that NetBackup supports

The cloud configuration files that are backed up during the catalog backup process reside at the following locations:

Windows `install_path\Veritas\NetBackup\var\global\wmc\cloud`

UNIX `/usr/opensv/var/global/wmc/cloud`

The files `CloudProvider.xml` and `cacert.pem` are at the following location:

Windows `<installed-path>\NetBackup\var\global\cloud`

UNIX `/usr/opensv/var/global/cloud/`

Note: The `cacert.pem` file is not backed up during the NetBackup catalog backup process.

This `cacert.pem` file is a cloud provider-specific file. This file is installed as part of the NetBackup installation. This file includes the well-known public cloud vendor CA certificates used by NetBackup.

Catalog backups

Because the catalog plays an integral part in a NetBackup environment, a special type of backup protects the catalog and is separate from regular client backups. A catalog backup policy backs up catalog-specific data as well as produces disaster recovery information. The catalog can be stored on a variety of media.

The catalog backup is designed for active environments in which continual backup activity occurs. It includes all the necessary catalog files, the databases (NBDB, NBAZDB, and BMRDB), and any catalog configuration files. The catalog backup can be performed while regular backup activity occurs. Incremental backups of a large catalog can significantly reduce backup times.

Configure a catalog backup before you run any regular backups. NetBackup needs information from the catalog to determine where the backups of files are located. Without a catalog, NetBackup cannot restore data.

See [“Configuring catalog backups”](#) on page 867.

As additional protection for the catalog, consider archiving the catalog.

From a catalog backup an administrator can recover either the entire catalog or pieces of the catalog. (For example, separately recover the databases from the configuration files.) Details about catalog recovery scenarios and procedures are available in the *NetBackup Troubleshooting Guide*.

The catalog backup process

The catalog backup performs the following tasks:

- Backs up the catalog while continual client backups are in progress.
- Performs a full or an incremental catalog backup.

- Runs the scheduled catalog backups.
- Copies the databases to the staging directory and then backs up that directory.
- Creates the disaster recovery package.
- Catalog backups that are made to tape also include the following items:
 - Spans multiple tapes for a catalog backup.
 - Allows for a flexible pool of catalog tapes.
Catalog backups to tape use media from the **CatalogBackup** volume pool only.
 - Appends to existing data on tape.
 - When an online catalog backup is run, it generates three jobs: A parent job, a child job for NetBackup relational database tables, and a child job for catalog images and configuration data. The child jobs contain the actual backed up data. Consider both child jobs to duplicate, verify, or expire the backup.

Refer to the following topics for information on how to configure a catalog backup:

See [“Prerequisites for backing up the NetBackup catalog”](#) on page 867.

See [“Configuring catalog backups”](#) on page 867.

Strategies that ensure successful NetBackup catalog backups

Use the following strategies to ensure successful catalog backups:

- Use only the methods that are described in this chapter to back up the catalogs. These are the only methods that can track all relevant NetBackup activities and ensure consistency between the catalog files.
- Back up the catalogs often. If catalog backup files are lost, the changes that were made between the last catalog backup and the time of the disk crash are lost.
- If you back up your catalogs to disk, always back up to a different disk than where the catalog files reside. If you back up the catalog to the disk where the actual catalog resides, both catalog backups are lost if the backup disk fails. Recovering the catalog is much more difficult. Also, ensure that the disk has enough space for the catalogs. Backups to a full disk fail.

Note: If a catalog backup is on tape, the tape must be removed when the backup is finished or regular backups cannot proceed. NetBackup does not mix catalog and regular backups on the same tape.

Prerequisites for backing up the NetBackup catalog

The following prerequisites exist for a catalog backup:

- Set a passphrase for the disaster recovery package.
See [“Disaster recovery packages”](#) on page 873.
See [“Setting a passphrase to encrypt disaster recovery packages”](#) on page 875.
If the passphrase is not set, catalog backups fail.
- The primary server and the media server must both be at the same NetBackup version.
See the [NetBackup Installation Guide](#) for information about mixed version support.
- Catalog backups write only to media in the **CatalogBackup** volume pool. A storage device must be configured and media must be available in the **CatalogBackup** volume pool.
- The following requirements exist if the primary server is configured to use a non-privileged user (or service user) account. For more information on this type of account, refer to the [NetBackup Security and Encryption Guide](#).
 - The service user account must have the write access permissions on the disaster recovery (DR) path.
 - Configure the catalog policy with the credentials for the service account. (This setting is located on the **Disaster recovery** tab.)
 - You cannot use another user account, even if that account has the access to the DR path. The NetBackup Administrator must ensure that the service user can write to any network share without switching the context to another user.
On Windows, this requirement is not applicable if the DR path is a network share.

Configuring catalog backups

To protect the NetBackup catalog, you create a backup policy that is specific for catalog backups.

For information on how to configure catalog backups in Windows clustered environments, see the [NetBackup Clustered Primary Server Administrator's Guide](#).

To configure a catalog backup

- 1 Review the prerequisites for performing catalog backups.
See [“Prerequisites for backing up the NetBackup catalog”](#) on page 867.
- 2 Sign in to the NetBackup web UI.

- 3 Click **Protection > Policies**. Then click **Add**.
- 4 On the **Attributes** tab, complete the following entries:
 - Enter a unique **Policy name**.
See [“NetBackup naming conventions”](#) on page 1054.
 - For the **Policy type**, select **NBU-Catalog**.
 - **Policy storage**
For disk storage units, increase the **Maximum Concurrent Jobs** storage unit setting to ensure that the catalog backup can proceed during regular backup activity.
See [“Maximum concurrent jobs storage unit setting”](#) on page 563.

Note: If your installation contains media servers at various versions, you can select a specific media server for the destination **Policy storage**. Do not select **Any Available**.

- **Policy volume pool**
NetBackup automatically creates a **CatalogBackup** volume pool that is selected by default only for **NBU-Catalog** policy types.
- For other policy attribute descriptions, see the following topic:
See [“Policy Attributes tab”](#) on page 663.
- 5 On the **Schedules** tab, configure the schedules you want for the catalog backup.

See [“Concurrently running catalog backups with other backups”](#) on page 870.

See [“Catalog policy schedule considerations”](#) on page 870.

See [“Schedule Attributes tab”](#) on page 731.
- 6 Click the **Disaster recovery** tab.

The tab contains information regarding the location of data crucial to disaster recovery.
 - Provide the path where each disaster recovery image file can be saved on disk. Enter the **Network share username** and **Network share password**, if necessary.

It is recommended that you use a network share or a removable device.
Do not save the disaster recovery information to the local computer.

- 7 Select **Send disaster recovery email** and enter one or more email addresses for NetBackup administrators (separated by commas).

After every catalog backup, NetBackup sends disaster recovery information to the administrators that are indicated here.

Make sure that email notification is enabled in your environment.

See [“Disaster recovery emails and the disaster recovery files”](#) on page 872.

- 8 Add the policies that back up any critical data to the **Critical policies** list.

These policies are any that you consider crucial to the recovery of a site in the event of a disaster. The disaster recovery report includes a list of the media that is used for backups of critical policies. The report includes media only for incremental and full backup schedules, so any critical policies should use only incremental or full backup schedules.

- 9 Click **Create**.

Backing up NetBackup catalogs manually

Catalog backups typically run automatically per the **NBU-Catalog** policy. You can also manually start a catalog backup.

A manual catalog backup is useful in the following situations:

- To perform an emergency backup. For example, if the system is scheduled to be moved and you cannot wait for the next scheduled catalog backup.
- If there is only one standalone drive and the standalone drive is used for catalog backups. In this situation, automatic backups are not convenient. The catalog backup tape must be inserted before each catalog backup and removed when the backup is done. (The tape swap is necessary because NetBackup does not mix catalog and regular backups on the same tape.)

To perform a manual catalog backup

- 1 Sign in to the NetBackup web UI.
- 2 Click **Protection > Policies**.
- 3 Select the catalog backup policy that you want to run.
- 4 Click **Manual backup**.
See [“Performing manual backups”](#) on page 834.
- 5 (Optional) Select the schedule that you want to use.
- 6 Click **Backup**.

Concurrently running catalog backups with other backups

You can schedule catalog backups to run concurrently with other backup types for the primary server.

Make the following adjustments to ensure that the catalog backup can proceed while regular backup activity occurs:

- Set the **Maximum jobs per client** value to greater than one. The property is found in the Global attributes host properties for the primary server.
See [“Global attributes properties”](#) on page 109.
- Increase the **Maximum concurrent jobs** setting on the storage unit where the backups are sent.
See [“Maximum concurrent jobs storage unit setting”](#) on page 563.

See [“Determining whether or not a catalog backup succeeded”](#) on page 871.

See [“Strategies that ensure successful NetBackup catalog backups”](#) on page 866.

Catalog policy schedule considerations

When you work with catalog policy schedules, consider the following:

- Schedule the catalog backups to occur on a regular basis. Without regular catalog backups, you risk losing regular backups if there is a problem with the disk that contains the catalogs.
- The following backup types are supported:
 - Full
 - Differential incremental
This incremental schedule depends on a full schedule.
 - Cumulative incremental
- The least frequent schedule runs if many schedules are due at the same time.
- One catalog backup policy can contain multiple incremental schedules that are session-based:
 - If one is cumulative and the others are differential, the cumulative runs when the backup session ends.
 - If all are cumulative or all are differential, the first schedule that is found runs when the backup session ends.
- The queued scheduled catalog backup is skipped if a catalog backup job from the same policy is running.

- Session end means that no jobs are running. (This calculation does not include catalog backup jobs.)
- The Vault catalog backup is run whenever triggered from Vault, regardless of whether a catalog backup job is running from the same policy.

How catalog incrementals and standard backups interact on UNIX

A catalog backup policy can include both full catalog backups and incremental catalog backups. However, incremental catalog backups differ from incremental standard backups. Catalog backups use both `mtime` and `ctime` to identify changed data. Standard incremental backups use only `mtime` to identify changed data.

Because of this difference, running a standard policy type backup that includes the `/usr/opensv/netbackup/db/images/` directory can adversely affect incremental catalog backups. When standard backups run, they reset the file access time (`atime`). In turn, the reset changes the `ctime` for files and directories. If an incremental catalog backup runs, it sees that the `ctime` has changed and backs up the files. The backup may be unnecessary since the files may not have changed since the last catalog backup.

To avoid additional processing during catalog backups, the following is recommended:

If incremental catalog backups are configured, exclude the NetBackup `/usr/opensv/netbackup/db/images/` directory from standard backups.

To exclude that directory, create a `/usr/opensv/netbackup/exclude_list` file on the primary server.

Determining whether or not a catalog backup succeeded

The All Log Entries, Problems, and Media Log reports, available from the Reports utility, provide information on NetBackup catalog backups. The All logs entries report is also available in the NetBackup web UI.

An email message is sent to the address that is indicated in the **Disaster recovery** settings for a catalog backup.

Configure this email with the `mail_dr_info.cmd` (on Windows) or the `mail_dr_info` script (on UNIX).

See the [Administrator's Guide, Volume II](#) for more information on setting up this script.

Recovering the catalog

Catalog recovery is discussed in the [NetBackup Troubleshooting Guide](#).

Disaster recovery emails and the disaster recovery files

In a catalog backup policy, you can configure the policy to send the disaster recovery information to an email address. This information appears on the **Disaster recovery** tab.

The disaster recovery email and the accompanying attachments that are sent contain the following important items for a successful catalog recovery:

- A list of the media that contains the catalog backup.
- A list of critical policies.
- Instructions for recovering the catalog.
- The image file as an attachment.

If a catalog backup policy included both full backups and incremental backups, the attached image file can be a full or an incremental catalog backup.

Recovering from an incremental catalog backup completely recovers the entire catalog if the **Automatically recover the entire NetBackup catalog** option is selected on the wizard panel. The entire catalog is recovered because the incremental catalog backup references information from the last full backup. You do not need to recover the last full catalog backup before you recover the subsequent incremental backups.

- The disaster recovery package (`.drpkg` file) as an attachment.

Note: If you are not able to receive the disaster recovery packages over emails even after the disaster recovery email configuration, and then ensure the following:

Your email exchange server is configured to have the attachment size equal to or greater than the disaster recovery package size. You can check the size of the package (`.drpkg` file size) on the disaster recovery file location that you have specified in the catalog backup policy.

The firewall and the antivirus software in your environment allows the files with the `.drpkg` extension (which is the extension of a disaster recovery package file).

NetBackup emails the disaster recovery file when the following events occur:

- The catalog is backed up.
- A catalog backup is duplicated or replicated.
- The primary catalog backup or any copy expires automatically or is expired manually.

On Windows: You can tailor the disaster recovery email process by providing the `mail_dr_info.cmd` script in the `install_path\Veritas\NetBackup\bin` directory. This script is similar to the `nbmail.cmd` script. See the comments in the `nbmail.cmd` script for use instructions.

Disaster recovery packages

For increased security, a disaster recovery package is created during each catalog backup. The disaster recovery package file has `.drpkg` extension.

The disaster recovery (DR) package stores the identity of the primary server host. NetBackup requires this package to get the identity of the primary server back after a disaster. Once you have recovered the host identity, you can perform the catalog recovery.

The disaster recovery package contains the following information:

- NetBackup CA-signed certificates and private keys of the primary server certificate and the NetBackup certificate authority (CA) certificate
- Information about the hosts in the domain
- Security settings
- External CA-signed certificates
External CA-signed certificates from Windows certificate store, if applicable
- NetBackup configuration options that are specific to external CA-signed certificates
- Key management service (KMS) configuration

Note: By default, the KMS configuration is not backed up during catalog backup. Set the `KMS_CONFIG_IN_CATALOG_BKUP` configuration option to 1 to include the KMS configuration as part of the disaster recovery package during catalog backup.

Note: You must set a passphrase for the disaster recovery package for the catalog backups to be successful.

About disaster recovery settings

For increased security, a disaster recovery package is created during each catalog backup.

See [“Disaster recovery packages”](#) on page 873.

During each catalog backup, a disaster recovery package is created and encrypted with the passphrase that you set. You need to provide this encryption passphrase while you install NetBackup on the primary server in a disaster recovery mode after a disaster.

The following options are displayed on the **Disaster recovery** tab:

Table 22-1 Disaster recovery settings

Setting	Description
Enter passphrase	Enter the passphrase to encrypt disaster recovery packages. - By default, the passphrase must contain a minimum of 8 and a maximum of 1024 characters. You can set the passphrase constraints using the <code>nbseccmd -setpassphraseconstraints</code> command option. - The existing passphrase and the new passphrase must be different. - Only the following characters are supported for the passphrase: White spaces, uppercase characters (A to Z), lowercase characters (a to z), numbers (0 to 9), and special characters. Special characters include: ~ ! @ # \$ % ^ & * () _ + - = ` { } [] : ; ' , . / ? < > "
Confirm passphrase	Re-enter the passphrase for confirmation.

Caution: Ensure that the passphrase contains only the supported characters. If you enter a character that is not supported, you may face problems during disaster recovery package restore. The passphrase may not be validated and you may not be able to restore the disaster recovery package.

Note the following before you modify the passphrase for the disaster recovery packages:

- Subsequent disaster recovery packages are encrypted with the new passphrase that you set.
- If you change the passphrase anytime, it is not changed for the previous disaster recovery packages. Only new disaster recovery packages are associated with the new passphrase.
- The passphrase that you provide when you install NetBackup on the primary server in the disaster recovery mode after a disaster must correspond to the disaster recovery package from which you want to recover the primary server host identity.

Setting a passphrase to encrypt disaster recovery packages

During each catalog backup, a disaster recovery package is created and encrypted with the passphrase that you set.

See [“Disaster recovery packages”](#) on page 873.

Workflow to set a passphrase to encrypt disaster recovery packages and use it after a disaster:

Review the following workflow to learn about disaster recovery package restore:

1. Set an encryption passphrase for disaster recovery packages.
2. Create a catalog policy.

See [“Configuring catalog backups”](#) on page 867.

Consider the following scenarios:

- If you have not set the passphrase earlier, NetBackup prevents you from configuring a new catalog backup policy.
- If the catalog backup policy is upgraded from a previous version, catalog backups continue to fail until the passphrase is set.

Note: Catalog backups may fail with status code 144 even though the passphrase is set. This issue occurs because the passphrase may be corrupted. To resolve this issue, you must reset the passphrase.

3. After a disaster, when you install NetBackup on the primary server in a disaster recovery mode, provide the passphrase that you have set earlier. NetBackup decrypts the disaster recovery package using this passphrase and gets the identity of the primary server back during installation.

Caution: Special action is required if you fail to provide the appropriate passphrase when you install NetBackup on the primary server after a disaster. You may need to redeploy the security certificates on all NetBackup hosts. For more details, refer to the following article:

https://www.veritas.com/content/support/en_US/article.100033743

4. Once the primary server identity is back in place, the secure communication between the primary server and the media server is established and you can perform catalog recovery.
5. After successful catalog recovery, you must set the disaster recovery package passphrase again, because the passphrase is not recovered during the catalog recovery. Catalog backups that you configure in a new NetBackup instance continue to fail until you set the passphrase.

To set or modify a passphrase

- 1 Open the NetBackup web UI.
- 2 At the top right, select **Settings > Global security**.
- 3 Select **Disaster recovery**.
See “[About disaster recovery settings](#)” on page 874.
- 4 Enter and confirm a passphrase.

Review the following password rules:

- The existing passphrase and the new passphrase must be different.
- By default, the passphrase must contain a minimum of 8 and a maximum of 1024 characters.

You can set the passphrase constraints using the `nbseccmd -setpassphraseconstraints` command option.

- Only the following characters are supported for the passphrase: White spaces, uppercase characters (A to Z), lowercase characters (a to z), numbers (0 to 9), and special characters. Special characters include: ~ ! @ # \$ % ^ & * () _ + - = ` { } [] | : ; ' , . / ? < > "

Caution: If you enter a character that is not supported, you may face problems during disaster recovery package restore. The passphrase may not be validated and you may not be able to restore the disaster recovery package.

- 5 Select **Save**. If the passphrase already exists, it is overwritten.

To set or modify a passphrase using the command-line interface

- 1 The NetBackup administrator must be logged on to the NetBackup Web Management Service to perform this task. Use the following command to log on:

```
bpnbat -login -loginType WEB
```

- 2 Run the following command to set a passphrase to encrypt disaster recovery packages:

```
nbseccmd -drpkgpassphrase
```

- 3 Enter the passphrase.

If a passphrase already exists, it is overwritten.

Archiving the catalog and restoring from the catalog archive

Catalog archiving helps administrators solve the kinds of problems that large amounts of catalog data can pose: large catalogs require a greater amount of disk space and can be time-consuming to back up.

Catalog archiving reduces the size of online catalog data by relocating the large catalog `.f` files to secondary storage. NetBackup administration continues to require regularly scheduled catalog backups, but the backups are faster without the large amount of online catalog data.

You can also use intelligent catalog archiving (ICA) to reduce the number of catalog `.f` files from secondary storage. When you enable ICA, any catalog `.f` file that is older than the specified retention period value is removed from the catalog disk. You can also specify a size value so that any catalog `.f` file that is greater than or equal to the size value is removed from the catalog disk.

See [“Enabling intelligent catalog archiving \(ICA\) to reduce the number of .f files”](#) on page 880.

Catalog archiving should not be used as a method to reclaim disk space when a catalog file system fills up. In that situation, investigate catalog compression or add disk space to grow the file system.

For additional catalog archiving considerations, see the following topic:

See [“Catalog archiving considerations”](#) on page 887.

To archive the catalog and restore the catalog archive

- 1 Use `bpcatlist` to determine what images are available to be archived.

Running `bpcatlist` alone does not modify any catalog images. Only when the `bpcatlist` output is piped to `bpcatarc` are the `.f` files backed up, and only when the output is piped to `bpcatrm` will the `.f` files be deleted from disk.

To determine what images have `.f` files on disk that can be archived, run the following command. The `catarcid` column indicates whether the `.f` file is not currently backed up (0) or the `catarcid` of the backup of that image.

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -online
```

To determine what images have been previously archived and removed from disk, run the following command.

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -offline
```

The catalog commands are described in detail in the following topic:

See [“Catalog archiving commands”](#) on page 885.

Note: If catalog archiving has not been previously run, this command should return: `No entity was found.`

For example, to display all images for a specific client before January 1, 2017, run the following command:

```
bpcatlist -client name -before Jan 1 2017
```

To display the help for the `bpcatlist` command run this command.

```
bpcatlist -help
```

Once the `bpcatlist` output correctly lists all the images that are to be archived or deleted, other commands can be added.

2 Running the catalog archive.

Before running the catalog archive, create a backup policy named **catarc**. The policy is required for the `bpcatarc` command to successfully process images. The name of the policy reflects that the purpose of the schedule is for catalog archiving.

See the following topic for details about configuring the **catarc** policy:

See [“Creating a catalog archiving policy”](#) on page 884.

To run the catalog archive, first run the `bpcatlist` command with the same options used in step 1 to display images. Then pipe the output through `bpcatarc` and `bpcatrm`.

```
bpcatlist -client all -before Jan 1 2017 | bpcatarc | bpcatrm
```

A new job appears in the **Activity Monitor**. The command waits until the backup completes before it returns the prompt. The command reports an error only if the catalog archive fails, otherwise the commands return to the prompt.

The **File List**: section of the Job Details in the **Activity Monitor** displays a list of image files that have been processed. When the job completes with a status 0, the `bpcatrm` command removes the corresponding `.f` files. If the job fails, no catalog `.f` files are removed.

If `bpcatlist` is piped to `bpcatarc` but the results are not piped to `bpcatrm`, the backup occurs but the `.f` files are not removed from disk. The same `bpcatlist` command can then be rerun and piped to `bpcatrm` to remove the `.f` files.

3 Restoring the catalog archive.

To restore the catalog archive, first use the `bpcatlist` command to list the files that need to be restored. Once `bpcatlist` displays the proper files to restore, run the `bpcatres` command to restore the actual files.

To restore all the archived files from step 2, run the following command:

```
bpcatlist -client all -before Jan 1 2017 | bpcatres
```

This command restores all of the catalog archive files before January 1, 2017.

Enabling intelligent catalog archiving (ICA) to reduce the number of .f files

Note: Intelligent catalog archiving (ICA) applies only to servers running NetBackup 11.1 and later using MSDP storage.

You can use intelligent catalog archiving (ICA) to reduce the number of catalog .*tf* files based on a specified retention period or file size. When you enable ICA, any catalog .*tf* file that is older than the specified retention period value is removed from the catalog disk. You can also specify a file size value so that any catalog .*tf* file that is greater than or equal to the size value is removed from the catalog disk.

The main advantage of ICA is that it shortens catalog backup time by reducing the number of .*tf* files that need to be backed up if they meet the required criteria:

- The backup image must be older than the configured ICA retention period.
- The .*tf* file must be larger than or equal to the configured ICA minimum size.
- At least one copy of the backup image must be on MSDP storage and have 1 or more true image restore (TIR) fragments.
- Image catalog .*tf* file has not been recalled in the last 24 hours.
- The backup image must be from a completed SLP or from a backup that is not managed by an SLP.
- The backup image is not from a catalog backup.
- The image catalog is not archived.

When ICA is enabled, you may notice the following behaviors:

- Initial image cleanup after you enable ICA may take longer than usual.
- Catalog backups will be faster if any of the .*tf* files that are involved have been intelligently archived.
- Browse and Restore functions will take longer if any of the .*tf* files that are involved have been intelligently archived.

No additional action is needed to restore the catalog .*tf* file. Catalog .*tf* files are restored from images automatically as follows:

- When an ICA image is browsed.
- When an ICA-eligible copy is expired from an ICA image. Restoring catalog .*tf* files ensures that the remaining copies from that image are accessible and usable.
- When an ICA-eligible image is found but its catalog .*tf* file missing.

More information about .*tf* files is available:

See [“About NetBackup image .*tf* files”](#) on page 863.

To enable intelligent catalog archiving (ICA) and specify retention and file size values

- 1 Run the following command on the primary server:

```
bpconfig -ica_retention seconds
```

When the *seconds* value is between 1 and 2147472000, ICA is enabled. Any image which is older than the value is processed for ICA. The catalog `.f` file from the ICA-eligible image is removed from the catalog disk. Setting this value to 0 (zero) disables ICA. The default value for NetBackup Flex Scale and Cloud Scale environments is 2592000 (30 days). The default value for all other NetBackup environments is 0 (disabled).

For Accelerator-enabled backups, specify an ICA retention value that is longer than full backup schedules so that the number of `.f` file restores from ICA images goes down.

For example, to set the ICA retention value to 30 days, enter `bpconfig -ica_retention 2592000`.

Use `bpconfig -U` to verify the change:

```
# bpconfig -U
Admin Mail Address:          sasquatch@wapati.edu
Job Retry Delay:             10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:                1 time(s) in 12 hour(s)
Keep Error/Debug Logs:      3 days
Max drives this master:      0
Keep TrueImageRecovery Info: 24 days
Compress DB Files:           (not enabled)
Media Mount Timeout:         30 minutes
Display Reports:             24 hours ago
Preprocess Interval:         0 hours
Image DB Cleanup Interval:    12 hours
Image DB Cleanup Wait Time:   10 minutes
Policy Update Interval:       10 minutes
Intelligent Catalog Archiving: Files file larger than 1024 KB
Intelligent Catalog Archiving: Images older than 30 day(s)
```

- 2 **Note:** After you enable ICA, the minimum file size for .f files is set to the default value 1024 KB. Use this step to change that value.

To specify a minimum file size, run the following command on the primary server:

```
bpconfig -ica_min_size size
```

When the *size* value is between 0 and 2097151, any catalog .f file that is larger than or equal to the size value is removed from the catalog disk. The default value is 1024.

For example to set the ICA minimum file size to 2048 KB, enter `bpconfig -ica_min_size 2048`.

Use `bpconfig -U` to verify the change:

```
# bpconfig -U
Admin Mail Address:          sasquatch@wapati.edu
Job Retry Delay:             10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:                1 time(s) in 12 hour(s)
Keep Error/Debug Logs:      3 days
Max drives this master:      0
Keep TrueImageRecovery Info: 24 days
Compress DB Files:           (not enabled)
Media Mount Timeout:         30 minutes
Display Reports:             24 hours ago
Preprocess Interval:         0 hours
Image DB Cleanup Interval:   12 hours
Image DB Cleanup Wait Time:  10 minutes
Policy Update Interval:      10 minutes
Intelligent Catalog Archiving: Files file larger than 2048 KB
Intelligent Catalog Archiving: Images older than 30 day(s)
```

To disable intelligent catalog archiving (ICA)

- ◆ Run the following command on the primary server:

```
bpconfig -ica_retention 0
```

Use `bpconfig -U` to verify the change:

```
# bpconfig -U
Admin Mail Address:          sasquatch@wapati.edu
Job Retry Delay:             10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:               1 time(s) in 12 hour(s)
Keep Error/Debug Logs:      3 days
Max drives this master:      0
Keep TrueImageRecovery Info: 24 days
Compress DB Files:           (not enabled)
Media Mount Timeout:         30 minutes
Display Reports:             24 hours ago
Preprocess Interval:         0 hours
Image DB Cleanup Interval:   12 hours
Image DB Cleanup Wait Time:  10 minutes
Policy Update Interval:      10 minutes
Intelligent Catalog Archiving: (not enabled)
```

Creating a catalog archiving policy

The catalog archiving feature requires the presence of a policy named **catarc** before the catalog archiving commands can run properly. The policy can be reused for catalog archiving.

To create a catalog archiving policy

- 1 Open the NetBackup web UI.
- 2 On the left, click **Protection > Policies**. Then click **Add**.
- 3 Enter the **Policy name catarc**.

The **catarc** policy waits until `bpcatarc` can activate it. Users do not run this policy. Instead, `bpcatarc` activates this special policy to perform a catalog backup job, then deactivates the policy after the job is done.

- 4 In the **Attributes** policy tab, set the **Policy type** to **Standard** or **MS-Windows**, according to the platform of the primary server.

- 5 In the **Attributes** policy tab, deactivate the catalog archive policy by clearing the **Go into effect at** box.
See “[Go into effect at \(policy attribute\)](#)” on page 681.
- 6 Select the **Schedules** tab and click **Add** to create a schedule.
In the **Attributes** schedule tab, the **Name** of the schedule is not restricted, but the **Type of backup** must be **User backup**.
- 7 Select a **Retention** for the catalog archive. Set the retention level for a time at least as long as the longest retention period of the backups being archived.
Data can be lost if the retention level of the catalog archive is not long enough.
You may find it useful to set up and then designate a special retention level for catalog archive images.
- 8 Select the **Start window** tab and define a schedule for the **catarc** policy.
The schedule must include in its window the time when the `bpcatarc` command is run. If the `bpcatarc` command is run outside of the schedule, the operation fails.
- 9 Click **Add** to save the schedule.
- 10 On the **Clients** tab, enter the name of the primary server as it appears on the NetBackup servers list.
- 11 On the **Backup selections** tab, browse to the directory where catalog backup images are placed:
On Windows: `install_path\NetBackup\db\images`
On UNIX: `/usr/openv/netbackup/db/images`
- 12 Click **Create** to save the policy.

Catalog archiving commands

The catalog archiving option relies on three commands to designate a list of catalog `.f` files, then archive the files. A fourth command, `bpcatres`, is used to restore the files if necessary.

Catalog archiving uses the following commands.

Table 22-2 Catalog archiving commands

Command	Description
bpcatlist	The <code>bpcatlist</code> command queries the catalog data. Then, <code>bpcatlist</code> lists the portions of the catalog that are based on selected parameters. For example, date, client, policy, schedule name, backup ID, the age of the backup image, or the date range of the backup image. <code>bpcatlist</code> outputs the formatted image summary information of matched images to standard output. The other catalog archiving commands, <code>bpcatarc</code>, <code>bpcatrm</code>, and <code>bpcatres</code>, all depend on input from <code>bpcatlist</code> by a piped command. For example, to archive (backup and delete) all of the <code>.f</code> files that were created before January 1, 2012, the following would be entered: <pre>bpcatlist -client all -before Jan 1 2012 bpcatarc bpcatrm</pre> <code>bpcatlist</code> is also used to provide status information. For each catalog, it lists the following information: ■ Backup ID (Backupid) ■ Backup date (Backup Date) ■ Catalog archive ID (catarcid). After one <code>.f</code> file is successfully backed up, a catalog archive ID is entered into the catarcid field in the image file. This field is zero (0) if the image was never archived. ■ Archived status (S). Indicates whether the catalog was archived (2) or was not archived (1). ■ Compressed status (C). Indicates whether the catalog was compressed (<i>positive_value</i>) or was not compressed (0). ■ Catalog file name (Files file) The following is an example of the <code>bpcatlist</code> output, showing all of the backups for client alpha since October 23: <pre># bpcatlist -client alpha -since Oct 23 Backupid Backup Date ...Catarcid S C Files file alpha_097238 Oct 24 10:47:12 2012 ... 973187218 1 0 alpha_097238_UBAK.f alpha_097233 Oct 23 22:32:56 2012 ... 973187218 1 0 alpha_097233_FULL.f alpha_097232 Oct 23 19:53:17 2012 ... 973187218 1 0 alpha_097232_UBAK.f</pre> More information is available in the NetBackup Commands Reference Guide.
bpcatarc	The <code>bpcatarc</code> command reads the output from <code>bpcatlist</code> and backs up the selected list of <code>.f</code> files. After one <code>.f</code> file is successfully backed up, a catalog archive ID is entered into the catarcid field in the image file. For archiving of the <code>.f</code> files to proceed, a policy by the name of catarc is required. The policy is based on a User Backup type schedule. The schedule for catarc must include in its window the time <code>bpcatarc</code> command is run. See “Creating a catalog archiving policy” on page 884.

Table 22-2 Catalog archiving commands (*continued*)

Command	Description
<code>bpcatrm</code>	The <code>bpcatrm</code> command reads the output from <code>bpcatlist</code> or <code>bpcatarc</code>. If the image file has valid catarcid entries, <code>bpcatrm</code> deletes selected image .f files from the online catalog. <code>bpcatrm</code> does not remove one .f file unless the file has been previously backed up using the catarc policy.
<code>bpcatres</code>	Use the <code>bpcatres</code> command to restore the catalog. The <code>bpcatres</code> command reads the output from <code>bpcatlist</code> and restores selected archived .f files to the catalog. For example: <pre>bpcatlist -client all -before Jan 1 2012 bpcatres</pre>

Catalog archiving considerations

Consider the following items before catalog archiving:

- Perform catalog archiving operations when NetBackup is in an inactive state (no jobs are running).
- Catalog archiving modifies existing catalog images. As a result, it should never be run when the catalog file system is 100% full.
- To ensure that catalog backup images are not on the same tapes as user backups, create a separate media pool for catalog archives.
- You may find it useful to set up and then designate, a special retention level for catalog archive images.
 To specify retention levels, open the NetBackup web UI. On the left click **Hosts > Host properties**. Locate the primary server and click **Edit primary server**. Then click **Retention periods**.
 See [“Retention periods properties”](#) on page 152.
- Additional time is required to mount the tape and perform the restore of archived .f files.
- There is no simple method to determine to which tape the catalog has been archived. The `bpcatlist -offline` command is the only administrative command to determine what images have been archived. This command does not list what tape was used for the archive. As a result, exercise caution to ensure that the tapes used for catalog archiving are available for restoring the archived catalog images. Either create a separate volume pool to use exclusively for catalog archives or find a method to label the tape as a catalog archive tape.

Extracting images from the catalog archives

The situation may arise in which a storage provider needs to extract all of a specific client's records. The storage provider can extract the customer images from the catalog archive by creating the archives that are based on client name.

To extract images from the catalog archives based on a specific client

- 1 Create a volume pool for the client.
- 2 Create a catalog archiving policy. Indicate the volume pool for that client in the **Attributes** tab.
- 3 Run `bpcatlist` so only the `.f` files from that client are listed. For example:

```
bpcatlist -client clientname | bpcatarc | bpcatrm
```
- 4 If you do not want to write more images to the client's volume pool, change the volume pool before you run another archiving catalog.

Estimating catalog space requirements

NetBackup requires disk space to store its error logs and information about the files it backs up.

The disk space that NetBackup needs varies according to the following factors:

- Number of files to be backed up
- Frequency of full and incremental backups
- Number of user backups and archives
- Retention period of backups
- Average length of full path of files
- File information (such as owner permissions)
- Average amount of error log information existing at any given time
- Whether you have enabled the database compression option.

To estimate the disk space that is required for a catalog backup

- 1 Estimate the maximum number of files that each schedule for each policy backs up during a single backup of all its clients.
- 2 Determine the frequency and the retention period of the full and the incremental backups for each policy.

- 3 Use the information from steps 1 and 2 to calculate the maximum number of files that exist at any given time.

For example:

Assume that you schedule full backups to occur every seven days. The full backups have a retention period of four weeks. Differential incremental backups are scheduled to run daily and have a retention period of one week.

The number of file paths you must allow space for is four times the number of files in a full backup. Add to that number one week's worth of incremental backups.

The following formula expresses the maximum number of files that can exist for each type of backup (daily or weekly, for example):

Files per Backup × Backups per Retention Period = Max Files

For example:

A daily differential incremental schedule backs up 1200 files and the retention period for the backup is seven days. Given this information, the maximum number of files that can exist at one time are the following:

$$1200 \times 7 \text{ days} = 8400$$

A weekly full backup schedule backs up 3000 files. The retention period is four weeks. The maximum number of files that can exist at one time are the following:

$$3000 \times 4 \text{ weeks} = 12,000$$

Obtain the total for a server by adding the maximum files for all the schedules together. Add the separate totals to get the maximum number of files that can exist at one time. For example, 20,400.

For the policies that collect true image restore information, an incremental backup collects catalog information on all files (as if it were a full backup). This changes the calculation in the example: the incremental changes from $1200 \times 7 = 8400$ to $3000 \times 7 = 21,000$. After 12,000 is added for the full backups, the total for the two schedules is 33,000 rather than 20,400.

- 4 Obtain the number of bytes by multiplying the number of files by the average number of bytes per file record.

If you are unsure of the average number of bytes per file record, use 132. The results from the examples in step 3 yield:

$$(8400 \times 132) + (12,000 \times 132) = 2692800 \text{ bytes (or about 2630 kilobytes)}$$

- 5 Add between 10 megabytes to 15 megabytes to the total sum that was calculated in step 4. The additional megabytes account for the average space that is required for the error logs. Increase the value if you anticipate problems.
- 6 Allocate space so all the data remains in a single partition.

NetBackup file size considerations on UNIX systems

File system limitations on UNIX include the following:

- Some UNIX systems have a large file support flag. Turn on the flag to enable large file support.
- Set the file size limit for the root user account to unlimited to support large file support.

Moving the image catalog

An image catalog may become too large for its current location. Consider moving the image catalog to a file system or disk partition that contains more available space.

Notes about moving the image catalog

- NetBackup does not support saving the catalog to a remote NFS share. CIFS is supported on some SAN or NAS storage.
See [“Parts of the NetBackup catalog”](#) on page 859.
- NetBackup only supports moving the image catalog to a different file system or disk partition. It does not support moving the other subdirectories that make up the entire NetBackup catalog.
For example, on Windows, do not use the `ALTPATH` mechanism to move `install_path\NetBackup\db\error`.
For example, on UNIX, do not move `/usr/opensv/netbackup/db/error`. The catalog backup only follows the symbolic link when backing up the `/images` directory. So, if symbolic links are used for other parts of the NetBackup catalog, the files in those parts are not included in the catalog backup.
- The directory that is specified in the `ALTPATH` file is not automatically removed if NetBackup is uninstalled. If NetBackup is uninstalled, you must manually remove the contents of this directory.

Moving the image catalog between Windows hosts

To move the image catalog on Windows

- 1 Back up the NetBackup catalogs manually.

A backup of the catalogs ensures that you can recover image information in case something is accidentally lost during the move.

See [“Backing up NetBackup catalogs manually”](#) on page 869.

- 2 Check the **Jobs** tab in the **Activity monitor** and ensure that no backups or restores are running for the client.

If jobs are running, either wait for them to end or stop them by using the **Jobs** tab in the Activity monitor.

- 3 Use the **Daemons** tab in the **Activity monitor** to stop the Request Manager and the Database Manager daemons. These services are stopped to prevent jobs from starting. Do not modify the database while this procedure is performed.

- 4 Create a file named `ALTPATH` in the image catalog directory.

For example, if NetBackup is installed in the default location and the client name is *mars*, the path to the image catalog is:

```
C:\Program Files\Cohesity  
NetBackup\NetBackup\db\images\mars\ALTPATH
```

- 5 Create the directory to which you intend to move the image information. For example:

```
E:\NetBackup\alternate_db\images\client_name
```

- 6 On the first line of the `ALTPATH` file, specify the path to the directory where you intend to move the client's image information. For example:

```
E:\NetBackup\alternate_db\images\client_name
```

The path is the only entry in the `ALTPATH` file.

- 7 Move all files and directories (except the `ALTPATH` file) that are in the current client directory to the new directory.

For example, if the images are currently in

```
C:\Program Files\Cohesity NetBackup\NetBackup\db\images\mars
```

and the `ALTPATH` file specifies

```
E:\NetBackup\alternate_db\images\mars
```

then move all files and directories (except the `ALTPATH` file) to

```
E:\NetBackup\alternate_db\images\mars
```

- 8 Start the NetBackup Request Daemon, NetBackup Job Manager, and NetBackup Policy Execution manager in the **Daemons** tab.

Backups and restores can now resume for the client.

Moving the image catalog between UNIX hosts

To move the image catalog on UNIX

- 1 Check that no backups are in progress by running:

```
/usr/opensv/netbackup/bin/bpps
```

- 2 Stop `bprd` by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

- 3 Stop `bpdbm` by running:

```
/usr/opensv/netbackup/bin/bpdbm -terminate
```

- 4 Create the directory in the new file system. For example:

```
mkdir /disk3/netbackup/db/images
```

- 5 Move the image catalog to the new location in the other file system.

- 6 Create a symbolic link from `/usr/opensv/netbackup/db/images` to the new location in the other file system.

See [“NetBackup file size considerations on UNIX systems”](#) on page 890.

About image catalog compression

The image catalog contains information about all client backups. It is accessed any time a user lists or restores files. NetBackup lets you compress all portions of the catalog or only older portions of the catalog.

Control image catalog compression by setting the **Compress catalog interval** in the **Global attributes** host property. This interval indicates how old the backup

information must be before it is compressed. Specify the number of days to defer compression information, so users who restore files from recent backups are not affected. By default, **Compress catalog interval** is set to 0 and image compression is not enabled.

See [“Global attributes properties”](#) on page 109.

Note: Cohesity discourages manually compressing or decompressing the catalog backups with the `bpimage -[de]compress` command or any other method. Manually compressing or decompressing a catalog backup while any backup (regular or catalog) is running results in inconsistent image catalog entries. When users list and restore files, the results can be incorrect.

It does not make a difference to NetBackup if the backup session was successful. The operation occurs while NetBackup expires backups and before it runs the `session_notify` script and the backup of the NetBackup catalogs.

The time to perform compression depends on the server speed and the number and size of the files being compressed. Files are compressed serially, and temporary working space is required in the same partition.

When numerous compressed image catalog files must be processed, the backup session is extended until compression is complete. The additional backup time is especially noticeable the first time you perform the compression. To minimize the effect of the initial sessions, consider compressing the files in stages. For example, begin by compressing the records for the backups older than 120 days. Continue to reduce the number of days over a period of time until you reach a comfortable setting.

Compressing the image catalog accomplishes the following objectives:

- Reduces greatly the disk space that is consumed.
- Reduces the media that is required to back up the catalog.

The amount of space that is reclaimed varies with the types of backups you perform. Full backups result in a larger percentage of catalog compression than incremental backups. Normally, more data is duplicated in a catalog file for a full backup. Using catalog compression, a reduction of 80% is possible.

This reduction in disk space and media requirements is achieved at the expense of performance when a user lists or restores files. Since the information is uncompressed at each reference, performance degradation is in direct proportion to the number and size of compressed files that are referenced. If the restore requires numerous catalog files to be uncompressed, increase the **File browse timeout** value that is associated with list requests. (See the **Timeouts** host property for the client.)

Uncompressing the NetBackup catalog

You may find it necessary to temporarily uncompress all records that are associated with an individual client. Uncompress the records if you anticipate large or numerous restore requests, for example.

To uncompress the NetBackup catalog on Windows

- 1 Verify that the partition where the image catalog resides contains enough space to accommodate the uncompressed catalog.
See [“Estimating catalog space requirements”](#) on page 888.
- 2 Stop the NetBackup Request Daemon service, `bprd`.
- 3 Verify that the NetBackup Database Manager, `bpdbm`, is running.
- 4 In the NetBackup web UI, select **Hosts > Host properties**.
- 5 Select the primary server and click **Connect**. Then select the server and click **Edit primary server**.
- 6 Select **Global attributes**.

See [“Global attributes properties”](#) on page 109.

- 7 Clear the **Compress catalog interval** check box. Then click **Save**.
- 8 Open a command prompt. Change to the following directory:

```
install_path\Veritas\NetBackup\bin\admincmd
```

Run one of the followings commands.

To decompress the records for a specific client, enter:

```
bpimage -decompress -client_name
```

To decompress the records for all clients, enter:

```
bpimage -decompress -allclients
```

- 9 Restart the NetBackup Request Daemon (`bprd`).
- 10 Restore the files from the client.
- 11 Set the **Compress catalog interval** to its previous value.

The records that were uncompressed for this client are compressed after the next backup schedule.

To uncompress the NetBackup catalog on UNIX

- 1 Perform the following steps as root on the primary server to uncompress the NetBackup catalog.

Verify that the partition where the image catalog resides has enough space to uncompress the client's image records.

- 2 Stop the request daemon, `bprd`, by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

- 3 Make sure that `bpdbm` is running:

```
/usr/opensv/netbackup/bin/bpps
```

- 4 In the NetBackup web UI, select **Hosts > Host properties**.

- 5 Select the primary server and click **Connect**. Then select the server and click **Edit primary server**.

- 6 Select **Global attributes**.

See [“Global attributes properties”](#) on page 109.

- 7 Clear the **Compress catalog interval** check box. Then click **Save**.

- 8 Change your working directory to `/usr/opensv/netbackup/bin` and run the command:

```
admincmd/bpimage -decompress -client name
```

- 9 Restart the request daemon `bprd`. Run the following command:

```
/usr/opensv/netbackup/bin/initbprd
```

- 10 Restore the files from the client.

- 11 Set the **Compress catalog interval** to its previous value.

The records that were uncompresssed for this client are compressed after the next backup schedule.

About the NetBackup database

This chapter includes the following topics:

- [About the NetBackup database installation](#)
- [Post-installation tasks](#)
- [Using the NetBackup Database Administration utility on Windows](#)
- [Using the NetBackup Database Administration utility on UNIX](#)

About the NetBackup database installation

Generally, the implementation of the NetBackup database in the NetBackup catalog is transparent. The NetBackup primary server includes a private, non-shared database server for the NetBackup database (NBDB).

The same installation of the NetBackup database is used for the optionally-licensed product, Bare Metal Restore (BMR) and its associated database (BMRDB). The BMR database is created during the BMR installation process.

By default, the NetBackup database (NBDB) is installed on the primary server. The primary server is also the default location for the Enterprise Media Manager (EMM). Since EMM is the primary user of NBDB, the NetBackup database always resides on the same computer as the Enterprise Media Manager.

See [“About the Enterprise Media Manager \(EMM\)”](#) on page 861.

About NetBackup primary server installed directories and files

The NetBackup Scale-Out Relational Database is installed in the following directories.

Windows

`install_path\Veritas\NetBackupDB`

`install_path\Veritas\NetBackup\bin`

`install_path\Veritas\NetBackupDB\data\instance`

The databases are installed in the following subdirectories:

`install_path\Veritas\NetBackupDB\data\nbdb\`

`install_path\Veritas\NetBackupDB\data\nbazdb\`

`install_path\Veritas\NetBackupDB\data\bmrdb\` (if BMR is installed)

On UNIX

`/usr/opensv/db`

`/usr/opensv/var/global`

`/usr/opensv/db/data/instance/`

The databases are installed in the following subdirectories:

`/usr/opensv/db/data/nbdb/`

`/usr/opensv/db/data/nbazdb/`

`/usr/opensv/db/data/bmrdb/`

About the `bin` directory

The `bin` is located as follows:

`install_path\Veritas\NetBackup\bin`

Warning: Use these utilities and commands in this directory with caution.

Contains the utilities and binaries for running and administering NetBackup services. More information can be found in the *NetBackup Commands Reference Guide*.

For information on using the NetBackup Database Administration utility (`NbDbAdmin.exe` or `dbadm`), see the following topics:

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 907.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 915.

About the contents of the NetBackupDB and db directories

The following table describes the contents of the following directories.

On Windows: *install_path*\Veritas\NetBackupDB\

On UNIX: */usr/opensv/db/*

Table 23-1 NetBackupDB and db directory contents

Directory	Description
bin	Contains the utilities and commands for administering the NetBackup database service.
data	The default location of the NetBackup databases (NBDB, NBAZDB, and BMRDB) and certain configuration files.
lib	On UNIX: Contains all the shared libraries for the NetBackup Scale-Out Relational Database. The directory also includes ODBC libraries, used to connect to NBDB and BMRDB.
scripts	Warning: Do not edit the scripts that are located in this directory. Contains the scripts that are used to create the NetBackup database. It also contains the scripts that are used to create the EMM and other schemas.
share	Contains the PostgreSQL document and module files that are required by the NetBackup database server.
staging	Used as a temporary staging area during catalog backup and recovery.
WIN64	(Windows) Contains .dll files for the NetBackup Scale-Out Relational Database.

About the data directory

The following directory is the default location of the NetBackup database, NBDB:

On Windows: *install_path*\NetBackupDB\data

On UNIX: */usr/opensv/db/data*

The *\data* directory contains the following subdirectories and files:

- *bmrdb*
If BMR is installed, this directory contains the BMR database.
- *nldb*
The main NetBackup database, including EMM.
- *nbazdb*
The NetBackup Authorization database.
- *vxdbs.conf*

The file that contains the configuration information specific to the installation of the NetBackup database.

See “[vxdbms.conf](#)” on page 899.

- `nbdbinfo.dat`
A backup of the NetBackup DBA password.

vxdbms.conf

On Windows:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_DATABASE = NBDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_AZ_DATABASE = NBAZDB
VXDBMS_NB_STAGING = C:\Program Files\Cohesity NetBackup\NetBackupDB\staging
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATA = C:\Program Files\Cohesity NetBackup\NetBackupDB\data
VXDBMS_NB_PASSWORD = encrypted_password
AZ_DB_PASSWORD = encrypted_password
VXDBMS_POSTGRESQL_POOLER_ODBC_PORT = 13787
```

On UNIX:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATABASE = NBDB
VXDBMS_AZ_DATABASE = NBAZDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_NB_DATA = /usr/opensv/db/data
VXDBMS_NB_STAGING = /usr/opensv/db/staging
VXDBMS_NB_PASSWORD = encrypted_password
AZ_DB_PASSWORD = encrypted_password
VXDBMS_POSTGRESQL_POOLER_ODBC_PORT = 13787
```

The encrypted password that is used to log into the DBA accounts is stored in `vxdbms.conf`. These accounts include NBDB, NBAZDB, and BMRDB and other data accounts.

NetBackup configuration entry

The `VXDBMS_NB_DATA` registry entry (Windows) or the `bp.conf` entry (UNIX) is a required entry and is created upon installation. The entry indicates the path to the directory where the following are located: NetBackup database, authorization database, BMR database, and the `vxdbms.conf` file.

On Windows:

HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\
 Config\VXDBMS_NB_DATA

On UNIX: /usr/openv/netbackup/bp.conf

VXDBMS_NB_DATA = /usr/openv/db/data

NetBackup database server management

This topic describes the commands that are available to manage the NetBackup database.

To start and stop the NetBackup database, use one of the following methods:

- In the **Daemons** tab of the Activity monitor, select the service **NetBackup Scale-Out Relational Database Manager** (vrtsdbsvc_psql).
- (Windows) From the Windows Service Manager, select the service **NetBackup Scale-Out Relational Database Manager** (vrtsdbsvc_psql).

- (Windows) Use the following commands:

```
install_path\Veritas\NetBackup\bin\bpdown -e vrtsdbsvc_psql
```

- `install_path\Veritas\NetBackup\bin\bpup -e vrtsdbsvc_psql`

- (UNIX) Use the following commands:

```
/usr/openv/db/bin/nbdbms_start_server -start
```

Starts the NetBackup Scale-Out Relational Database server if no option is specified.

```
/usr/openv/db/bin/nbdbms_start_server -stop -f
```

Stops the server; `-f` forces a shutdown with active connections.

The **NetBackup Scale-Out Relational Database Manager** daemon is included in the `stop` command or the `start` command, which starts and stops all NetBackup daemons.

Individual databases can be started or stopped, while the NetBackup Scale-Out Relational Database Manager service continues. Use the NetBackup Database Administration utility or the following commands:

- `nbdb_admin [-start | -stop]`

Starts or stops NBDB without shutting down the NetBackup Scale-Out Relational Database server.

To see whether the database is up, enter `nbdb_ping`.

- `nbdb_admin [-start | -stop BMRDB]`

Starts or stops BMRDB without shutting down the NetBackup Scale-Out Relational Database server.

To see whether the BMRDB database is up, enter `nbdb_ping -dbn BMRDB`.

The NetBackup database and clustered environments

The NetBackup database is supported in a clustered environment. Failover is included with the NetBackup server failover solution. The software is installed on all computers in the cluster.

The databases and the configuration files are installed in the following shared locations.

Windows

NetBackup databases:

`shared_drive\VERITAS\NetBackupDB\data`

Configuration files:

`shared_drive\VERITAS\NetBackupDB\data\instance`

UNIX

NetBackup databases:

`shared_drive/db/data`

Configuration files:

`/usr/opensv/var/global`

`shared_drive/db/data/instance`

Post-installation tasks

The tasks that are described in the following topics are optional and can be performed after the initial installation:

- Change the database password.
See [“Changing the NetBackup database password”](#) on page 902.
- Move the NetBackup databases (possibly to tune performance).
See [“Moving a database after installation ”](#) on page 903.
- Recreate NBDB.
See [“Creating the NBDB database manually”](#) on page 905.

Commands and utilities for administering the NetBackup databases

Note: Using the database administration utilities to administer the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Only use these utilities and commands with assistance of Cohesity Technical Support.

The following utilities are available to administer the databases.

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 907.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 915.

Also see the following commands in the *NetBackup Commands Reference Guide*.

`create_nbdb`

`nbdb_backup`

`nbdb_restore`

`nbdb_unload`

Changing the NetBackup database password

The database password is set to a randomly generated password upon installation. This password is used for NBDB and BMRDB and for all DBA and application accounts. You can use this procedure to change it to a known password.

The password is encrypted and stored in the `vxdbsms.conf` file. The permissions for the `vxdbsms.conf` file allow only a Windows administrator or a `root` user to read or write to it.

For requirements when NBAC is enabled, see the *NetBackup Security and Encryption Guide*.

To change the database password

- 1 Log on to the server as a Windows Administrator or as `root`.
- 2 To change the password for the first time after installation, run the following command. The command updates the `vxdbms.conf` file with the new, encrypted string:

On Windows: `install_path\NetBackup\bin\nbdb_admin -dba new_password`

On UNIX: `/usr/opensv/db/bin/nbdb_admin -dba new_password`

The password needs to be an ASCII string. Non-ASCII characters are not allowed in the password string.

- 3 To change a known password to a new password, you can either use the `nbdb_admin` command or the NetBackup Database Administration utility. You must know the current password to log into the NetBackup Database Administration utility.

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 907.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 915.

Moving a database after installation

The NetBackup database (NBDB) and the NetBackup authorization database (NBAZDB), are created on the primary server by default. To improve performance, you can use the NetBackup database administration utilities or command-line options to change the location of the databases.

Note the following:

- If BMR is installed and you want to move its database, it must reside on the primary server.
- Due to performance issues, you can only move a database to another disk or volume. The disk or volume must be locally attached.
NetBackup does not support saving the NetBackup database (NBDB, including EMM), NBAZDB, or the configuration files to a remote NFS share. CIFS is supported on some SAN storage and NAS storage.
- Run a catalog backup to back up NBDB and BMRDB both before and after moving the databases.

Moving a NetBackup database on Windows

The following instructions describe how to use the database administration utility to move a database.

You can also use the following command:

```
install_path\Veritas\NetBackup\bin\nbdb_move.exe
```

You can run the `nbdb_move` command at any time because it does not drop the database and recreate it. Therefore all the data is preserved.

To move a NetBackup database on Windows

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup services by typing the following command:

```
install_path\Veritas\NetBackup\bin\bpdown
```
- 3 Start the NetBackup Scale-Out Relational Database Manager service:

```
install_path\Veritas\NetBackup\bin\bpup -e vrtsdbsvc_psql
```
- 4 Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 5 From the **Database** list, select the database that you want to move.
- 6 Select the **Tools** tab.
- 7 Click **Move**.
- 8 Select **Move data to** and browse to the new location.
- 9 NetBackup does not require that the database directories are world-writable. Make sure that the new database directories (*data_directory*) have appropriate permissions so that the directories are not world-writable.
- 10 Start all services by typing the following command:

```
install_path\Veritas\NetBackup\bin\bpup
```
- 11 Perform a catalog backup.

Moving a NetBackup database on UNIX

To move a NetBackup database on UNIX

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```
- 3 Start the NetBackup Scale-Out Relational Database Manager daemon:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```
- 4 Use one of the following methods to move the existing databases:

- Use the **Move Database** option in the NetBackup Database Administration utility (dbadm).

- Enter the following command:

```
/usr/opensv/db/bin/nbdb_move  
-data data_directory
```

You can run the `nbdb_move` command at any time because it does not drop the database and recreate it. Thus, all data is preserved.

```
/usr/opensv/db/bin/nbdb_move -data data_directory
```

Note: NetBackup does not require that the database directories are world-writable. Make sure that the new database directories (`data_directory`) have appropriate permissions so that the directories are not world-writable.

- 5 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 6 Perform a catalog backup.

Copying the NetBackup databases

A temporary backup of the NBDB, NBAZDB, and BMRDB databases can be made for extra protection before database administration activities such as moving or reorganizing the databases. Also, some customer support situations may require that you create a copy of the NetBackup database.

Use the NetBackup database administration utilities or the `nbdb_backup` command to make this kind of backup.

Creating the NBDB database manually

The NBDB database is created automatically during NetBackup installation. However, it may be necessary during certain catalog recovery situations to create it manually by using the `create_nbdb` command.

Caution: Recreating the database manually is not recommended in most situations.

Note: If the NBDB database already exists, the `create_nbdb` command does not overwrite it. If you want to move the database, move it by using the `nbdb_move` command.

To create the NBDB database manually on Windows

- 1 Shut down all NetBackup services by typing the following command:

```
install_path\Veritas\NetBackup\bin\bpdown
```

- 2 Start the NetBackup Scale-Out Relational Database Manager service with the following command:

```
install_path\Veritas\NetBackup\bin\bpup -e vrtsdbsvc_psql
```

- 3 Run the following command:

```
install_path\Veritas\NetBackup\bin\create_nbdb.exe
```

- 4 Start all NetBackup services by typing the following command:

```
install_path\Veritas\NetBackup\bin\bpup
```

- 5 The new NBDB database is empty and does not contain the EMM data that is loaded during a normal installation.

Make sure that you have the most current support for new devices before the data is repopulated. New devices are added approximately every 2 months.

- 6 Repopulate the EMM data by running the `tpext` utility. `tpext` updates the NetBackup database with new versions of device mappings and external attribute files.

```
install_path\Veritas\Volmgr\bin\tpext.exe
```

During regular installation, `tpext` is run automatically.

If the `create_nbdb` command is used to create a database manually, the `tpext` utility must also be run. `tpext` loads EMM data into the database.

To create the NBDB database manually on UNIX

- 1 Shut down all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- 2 Start the NetBackup Scale-Out Relational Database Manager service with the following command:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

- 3 Run the following command:

```
/usr/opensv/db/bin/create_nbdb
```

- 4 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 5 The new `NBDB` database is empty and does not contain the `EMM` data that is loaded during a normal installation.

Make sure that you have the most current support for new devices before the data is repopulated. New devices are added approximately every 2 months.

- 6 Repopulate the `EMM` data by running the `tpext` utility. `tpext` updates the NetBackup database with new versions of device mappings and external attribute files.

```
/usr/opensv/volmgr/bin/tpext
```

During regular installation, `tpext` is run automatically.

If the `create_nbdb` command is used to create a database manually, the `tpext` utility must also be run. `tpext` loads `EMM` data into the database.

Additional `create_nbdb` options

In addition to using the `create_nbdb` command to create the `NBDB` database, you also can use it to perform the following actions. In each command, *NB_server_name* matches the name in the following file: `postgresql.conf`

- Drop the existing `NBDB` database and recreate it in the default location:

```
create_nbdb -drop
```

On UNIX, the location of the current `NBDB` data directory is retrieved automatically from the `bp.conf` file.

- Drop the existing `NBDB` database and do not recreate it:

```
create_nbdb -drop_only
```

- Drop the existing `NBDB` database and recreate it in the *data* directory:

```
create_nbdb -drop -data data_directory
```

If the `NBDB` database was moved from the default location by using `nbdb_move`, use this command to recreate it in the same location. Specify `current_data_directory`. `BMRDB` must also be recreated. The `BMRDB` database must reside in the same location as the NetBackup database.

Using the NetBackup Database Administration utility on Windows

The NetBackup administrator can use the Database Administration utility to configure the NetBackup databases and to monitor database operations. To use the utility, the administrator must have Administrator user privileges.

The NetBackup Database Administration utility is a standalone application (NbDbAdmin.exe) and is located in the following directory:

install_path\NetBackup\bin\NbDbAdmin.exe

To use the utility, you must be an administrator with administrator privileges.

When you start the NetBackup Database Administration utility, enter the DBA password. The password is set to a randomly generated password upon installation. Use the `nbdb_admin` command to change it to a known password if you have not done so already.

See [“Changing the NetBackup database password”](#) on page 902.

The NetBackup Database Administration utility displays the following information:

Table 23-2 NetBackup Database Administration properties

Property	Description
Database name and status	Select the database to administer. The list of possible databases is derived from the <code>vxdbms.conf</code> file. The <code>vxdbms.conf</code> file is located in the directory that is specified in the Windows registry parameter <code>VXDBMS_NB_DATA</code>. The database must reside on the same computer where the NetBackup Database Administration utility runs. One of the following status reports display for the selected database: ■ If the database is available, the screen displays Alive and well. ■ If the database is unavailable, the screen displays Not available.
Stop	Shuts down the selected database.
Start	Starts the selected database.
General tab	Contains the information about the database tablespaces. See “General tab of the NetBackup Database Administration utility” on page 909.
Tools tab	Contains a variety of tools to administer the selected database. See “Tools tab of the NetBackup Database Administration utility” on page 911.

Table 23-2 NetBackup Database Administration properties (*continued*)

Property	Description
Drive Space	Displays the amount of free space and used space on a drive. The Drive Space dialog displays the following information: ■ Drive ■ Capacity ■ Used space ■ Free space ■ % Utilized ■ Space

General tab of the NetBackup Database Administration utility

The **General** tab contains information about database tablespaces. The tab contains tools to let the administrator reorganize fragmented database objects and validate and rebuild the database.

Table 23-3 General tab options

Option	Description
Refresh	Displays the most current information.
Reorganize All	This option defragments the tablespaces that are fragmented.

Table 23-3 General tab options (*continued*)

Option	Description
Validate	This option performs a database validation on all of the database tablespaces in the selected database. - Validates the indexes and keys on all of the tables in the database. - Scans each table. For each row, a check is made that it exists in the appropriate indexes. The number of rows in the table must match the number of entries in the index. - Ensures that every row that is referenced in each index exists in the corresponding table. For foreign key indexes, it also ensures that the corresponding row exists in the primary table. After a validation check runs, the Results screen lists each database object. Each error is listed next to the database object where it was found. The total number of errors are listed at the end of the list of database objects. If no errors were found, that is indicated. If any validation errors are reported, perform the following tasks: - Shut down NetBackup (all daemons and services). - Start only the NetBackup database server (vrtssdbsvc_psql). - Click Validate to repeat the validation check or use the <code>nbdb_admin.exe</code> command line utility. If validation errors persist, contact Cohesity Technical Support. The administrator may be asked to rebuild the database using the Rebuild option or the <code>nbdb_unload.exe</code> command line utility.
Rebuild	This option unloads and reloads the database. A new database with all of the same options is built in its place. A Database Rebuild may be required if validation errors are reported when you use the Validate option. Note: Before you rebuild the database, it is recommended that you create a copy of the database by performing a backup from the Tools tab. To rebuild the database temporarily suspends NetBackup operations and can take a long time depending on the database size.

About fragmentation

Table fragmentation can impede performance. When rows are not stored contiguously, or if rows are split into more than one page, performance decreases because these rows require additional page accesses.

When an update to a row causes it to grow beyond the originally allocated space, the row is split. The initial row location contains a pointer to another page where the entire row is stored. As more rows are stored on separate pages, more time is required to access the additional pages.

Reorganizing may also reduce the total number of pages that are used to store the table and its indexes. It may reduce the number of levels in an index tree. Note that the reorganization does not result in a reduction of the total size of the database.

The **Rebuild** option on the **General** tab completely rebuilds the database, eliminating any fragmentation, and free space. This option may result in a reduction of the total size of the database.

See [“Estimating catalog space requirements”](#) on page 888.

Tools tab of the NetBackup Database Administration utility

The **Tools** tab of the NetBackup Database Administration utility contains a variety of tools to administer the selected database:

Password	See “Changing the DBA password using the NetBackup Database Administration utility” on page 911.
Move Database	See “Moving a NetBackup database” on page 912.
Unload	See “Exporting database schema and data” on page 912.
Backup	See “Copying or backing up a database” on page 913.
Restore	See “Restoring a database from a backup” on page 914.

Changing the DBA password using the NetBackup Database Administration utility

To log into the Database Administration utility, you must know the current DBA password.

To change the password for the first time after installation, use the `nbdb_admin` command. The command updates the `vxdbs.conf` file with the new, encrypted string:

See [“Changing the NetBackup database password”](#) on page 902.

To change a known password to a new password, you can either use the `nbdb_admin` command or the NetBackup Database Administration utility.

To change the DBA password from a known password to a new password

- 1 Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 2 Select the **Tools** tab.
- 3 In the **Password** section, click **Change**.

- 4** Enter the new password and confirm the new password. Changing the password changes it for both NBDB and BMRDB, if a BMR database is present.
- 5** Enable **Create a backup file of your new DBA password** to keep track of the password.
- 6** Click **OK**.

The utility warns you that it is important to remember the password. You cannot recover information within the NetBackup database if the password is unavailable.
- 7** Restart the database for the password change to take effect.

Moving a NetBackup database

Use the NetBackup Database Administration utility to change the location of a database.

For full instructions on how to move a database, see the following topic.

See [“Moving a database after installation ”](#) on page 903.

Exporting database schema and data

Use the NetBackup Database Administration utility to unload either the schema or the schema and data from the NetBackup database.

To export database schema and data

- 1** Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**
- 2** Select the **Tools** tab.
- 3** In the **Unload** section, click **Export**.
- 4** Browse to a destination directory.

5 Select one or more of the following options:

Schema	Unload only the database schema. The schema is unloaded as a file that is named <i>database.sql</i> in the named directory. For the NBDB database, the schema is unloaded as a file that is named <i>NBDB.sql</i> in the named directory. For other databases, a similar file is created. For example, for BMRDB the file is <i>BMRDB.sql</i> . For NBAZDB the file is <i>NBAZDB.sql</i> .
Schema and data	Unload both the database schema and the data. The data is unloaded as a set of files in comma-delimited format. One file is created for each database table.

6 Click **OK**.

Copying or backing up a database

Use the NetBackup Database Administration utility to back up the database to a specified directory.

It is recommended that you create a backup copy of a database in the following situations:

Before you move the database.	See “Moving a NetBackup database” on page 912.
Before you rebuild the database.	See “General tab of the NetBackup Database Administration utility” on page 909.

Note: Using the NetBackup Database Administration utility to back up and restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use the tool to back up and restore the NetBackup catalog only as a precautionary measure.

To copy or back up a database

- 1** Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 2** Select the **Tools** tab.
- 3** Click **Copy**.

- 4 Browse to a destination directory.

A copy of the database is made to this directory. This directory is also the location of the database that the **Restore** option uses.

Note: This backup is not a catalog backup, performed as part of regular NetBackup operations.

See [“Restoring a database from a backup”](#) on page 914.

- 5 Click **OK**.

Restoring a database from a backup

Use the NetBackup Database Administration utility to restore a database from a backup copy.

The restore overwrites the current database. The database is shut down and restarted after the restore is completed.

A database restore causes NetBackup activity to be suspended, so do not perform a database restore while active backups or other restores run.

Note: Using the Database Administration utility to back up and restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use the tool to back up and restore the NetBackup database only as a precautionary measure.

To restore a database from a backup

- 1 Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 2 Select the **Tools** tab.
- 3 Click **Restore**.
- 4 Browse to the directory that contains the backup database.
- 5 Click **OK**.

Using the NetBackup Database Administration utility on UNIX

The NetBackup Database Administration utility (`dbadm`) is a standalone application that is supported for NBDB and BMRDB. It is installed in the following location:

```
/usr/opensv/db/bin
```

To use the NetBackup Database Administration utility, you must be an administrator with root user privileges. When you start the NetBackup Database Administration utility, enter the DBA password. The password is set to a randomly generated password upon installation. Use the `nbdb_admin` command to change it to a known password if you have not done so already.

See [“Changing the NetBackup database password”](#) on page 902.

After you log on, the NetBackup Database Administration utility displays the following information about the current database:

Table 23-4 NetBackup Database Administration utility properties

Property	Description
Selected database	The selected database: NBDB or BMRDB
Status	The status of the selected database: UP or DOWN
Consistency	The validation state of the selected database: OK, NOT_OK, or DOWN

The initial screen also displays the following Database Administration main menu:

Table 23-5 Database Administration main menu options

Option	Description
Select/Restart Database and Change Password	This option displays the menu where you can select a database to start or stop, and to change database passwords. See “Select/Restart Database and Change Password menu options” on page 916.
Database Space Management	This option displays the menu where you can perform the following actions: ■ Generate a database space utilization report ■ Reorganize fragmented database objects See “Database Space Management menu options” on page 917.
Transaction Log Management	This option is not supported.

Table 23-5 Database Administration main menu options (*continued*)

Option	Description
Database Validation Check and Rebuild	This option displays the menu where you can validate and rebuild the selected database. See “Database Validation Check and Rebuild menu options” on page 918.
Move Database	This option displays the menu where you can change the location of the database tablespaces. See “Move Database menu options” on page 919.
Unload Database	This option displays the menu where you can unload either the schema or the schema and data from the database. See “Unload Database menu options” on page 920.
Backup and Restore Database	This option displays the menu where you can choose the backup and restore options for the database. See “Backup and Restore Database menu options” on page 920.
Refresh Database Status	This option refreshes the Status and Consistency in the main menu.

Select/Restart Database and Change Password menu options

The Select/Restart Database and Change Password menu contains the following options.

Table 23-6 Select/Restart Database and Change Password options

Option	Description
NBDB	Select NBDB and then view or modify the database using the other <code>dbadm</code> menu options.
BMRDB	Select BMRDB and then view or modify the database using the other <code>dbadm</code> menu options.
Start Selected Database	Starts the selected database.
Stop Selected Database	Stops the selected database.

Table 23-6 Select/Restart Database and Change Password options
(continued)

Option	Description
Change Password	Changes the password for the databases. The password is changed for both NBDB and BMRDB, if applicable. Restart the database for the password change to take effect. To log into the Database Administration utility, you must know the current DBA password. To change the password for the first time after installation, use the <code>nbdb_admin</code> command. The command updates the <code>vxdbms.conf</code> file with the new, encrypted string: See “Changing the NetBackup database password” on page 902. To change a known password to a new password, you can either use the <code>nbdb_admin</code> command or the NetBackup Database Administration utility.

Database Space Management menu options

You can use the Database Space Management option to perform the following functions:

- To report on database space utilization
- To reorganize fragmented database objects

Table 23-7 Database Space and Memory Management options

Option	Description
Report on Database Space	The report contains the tablespaces and the physical pathnames of the databases. For each tablespace, the report displays the name, the amount of free space in KBytes, and the file size in KBytes. The report also displays the amount of free space that remains on each of the file systems being used for the database.

Table 23-7 Database Space and Memory Management options (*continued*)

Option	Description
Database Reorganize	Select this option to reorganize fragmented database tablespaces. These actions are performed from the Database Reorganize menu as follows: ■ 1) Defragment All This option automatically determines the tablespaces that are fragmented. ■ 2) Table Level Defragmentation This option generates a fragmentation report for each database table. For each table, the report includes the TABLE_NAME, number of ROWS, number of ROW_SEGMENTS, and SEGS_PER_ROW. In addition, a * displays in the ! column for an individual table if it will be automatically selected for reorganization by the Defragment All option. A row segment is all or part of one row that is contained on one page. A row may have one or more row segments. The ROW_SEGMENTS value indicates total number of row segments for the table. The SEGS_PER_ROW value shows the average number of segments per row, and indicates whether or not a table is fragmented. A SEGS_PER_ROW value of 1 is ideal, and any value more than 1 indicates a high degree of fragmentation. For example, a value of 1.5 means that half of the rows are partitioned. See “About fragmentation” on page 910.

Database Validation Check and Rebuild menu options

The Database Validation Check and Rebuild option lets you validate and rebuild the currently selected database.

Table 23-8 Database Validation Check and Rebuild menu options

Option	Description
Standard Validation	The standard type of validation is not supported. This option performs a full validation.

Table 23-8 Database Validation Check and Rebuild menu options (*continued*)

Option	Description
Full Validation	This option performs a database validation on all of the database tablespaces in the selected database. - Validates the indexes and keys on all of the tables in the database. - Scans each table. For each row, a check is made that it exists in the appropriate indexes. The number of rows in the table must match the number of entries in the index. - Ensures that every row that is referenced in each index exists in the corresponding table. For foreign key indexes, it also ensures that the corresponding row exists in the primary table. Note: To perform a full database validation, shut down NetBackup and start only the database service. If any validation errors are reported, perform the following tasks: - Shut down NetBackup (all daemons and services). - Start only the NetBackup database server (vrtsdbsvc_psqli). - Repeat the validation check using this tool or the <code>nbdb_admin</code> command line utility. If validation errors persist, contact Cohesity Technical Support. The administrator may be asked to rebuild the database using the Database Rebuild option or the <code>nbdb_unload.exe</code> command-line utility.
Database Rebuild	This option lets you rebuild the database. A Database Rebuild results in a complete unload and reload of the database. A new database with all of the same options is built in place. A Database Rebuild may be required if Database Validation errors are reported using the Standard or Full Validation options. During a Database Rebuild, all NetBackup operations are suspended. When you select this option, a message appears which recommends that you exit and create a backup using the Backup Database option before you rebuild the database. You then have the choice of whether to continue or not. See “Backup and Restore Database menu options” on page 920.

Move Database menu options

The Move Database menu option lets you change the location of a database. After you select Move Database, you are prompted for the directory name where you want to move the database.

For full instructions on how to move a database, see the following topic.

See [“Moving a database after installation ”](#) on page 903.

Unload Database menu options

The Unload Database menu options let you unload either the schema or the schema and data from the `NBDB` or the `BMRDB` database.

A file is created that can be used to rebuild the database. If the data is also included in the unload, a set of data files in comma-delimited format is created.

The Unload Database menu contains the following options.

Table 23-9 Unload Database menu options

Option	Description
Schema Only	This option lets you unload only the database schema. For the <code>NBDB</code> database, the schema is unloaded as a file that is named <code>NBDB.sql</code> in the named directory. For <code>BMRDB</code> the file is <code>BMRDB.sql</code> .
Data and Schema	This option lets you unload both the database schema and the data. The data is unloaded as a set of files. One file is created for each database table.
Change Directory	This option lets you change the directory location for the files that unload options (1) or (2) create.

Backup and Restore Database menu options

The Backup and Restore Database menu options let you back up the NetBackup database to the specified directory. You can restore from a previously created backup.

It is recommended to create a backup copy of the databases in the following situations:

- Before you move the database.
- Before you rebuild the database.

Note: Using the NetBackup Database Administration utility to back up and restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use the tool to back up and restore the NetBackup database only as a precautionary measure.

Table 23-10 Backup and Restore Database menu options

Option	Description
Online Backup	This option lets you make a copy of the databases while the databases are active. Other NetBackup activity is not suspended during this time.
Restore Backup	This option lets you restore from a copy of the databases that was previously made with either options 1 or 2. The currently running databases are overwritten, and the database is shut down and restarted after the restore is completed.
Change Directory	This option lets you change the directory location for the databases that the backup options (1) or (2) create. This directory is the source of the databases for the restore option (3).

Managing backup images

This chapter includes the following topics:

- [About the Catalog utility](#)
- [Catalog utility search criteria and backup image details](#)
- [Verifying backup images](#)
- [Promoting a copy to a primary copy](#)
- [Duplicating backup images](#)
- [Expiring backup images](#)
- [About Image Dependency Expiration Cleanup](#)
- [About importing backup images](#)

About the Catalog utility

Use the **Catalog** utility to create and configure catalog backups. Catalog backups are required for NetBackup to protect NetBackup internal databases. The catalogs contain setup information as well as critical information about client backups. The catalog backups are tracked separately from other backups to ensure recovery in case of a server crash.

The **Catalog** utility is also used to perform the following actions:

- Search for backup images to verify the contents of media with what is recorded in the NetBackup catalog.
See [“Catalog utility search criteria and backup image details”](#) on page 923.
- Duplicate a backup image.
See [“Duplicating backup images”](#) on page 928.
- Promote a backup image from a copy to the primary backup copy.

- See [“Promoting a copy to a primary copy”](#) on page 927.
- Expire backup images.
 See [“Expiring backup images”](#) on page 933.
- Import expired backup images or images from another NetBackup server.
 See [“About importing expired images”](#) on page 935.

Catalog utility search criteria and backup image details

The catalog utility in the NetBackup web UI lets you perform various actions on a catalog image. For example, verify or duplicate an image. The catalog utility is organized as follows:

- **Search** tab
 Provides the search criteria you can use to locate backup images. See [Table 24-1](#) for details.
 For more details on these actions and on data-in-transit encryption (DTE) in your NetBackup environment, see the [NetBackup Administrator's Guide, Volume I](#) and [NetBackup Security and Encryption Guide](#).
 After you search for backup images, the image list displays at the bottom of the page. Click **Show or hide columns** to display additional information about the images. See [Search results properties](#) for additional properties that are displayed in the search results.
- **Activity** tab
 Displays the progress of the request to verify, duplicate, expire, or import an image.

Search criteria

The following actions and search criteria are available when you search for catalog images.

Table 24-1 Catalog search criteria

Property	Description
Action	Specifies the action that was used to create the image: Verify, Duplicate, Import. See “Verifying backup images” on page 926. See “Duplicating backup images” on page 928. See “Expiring backup images” on page 933.

Table 24-1 Catalog search criteria (*continued*)

Property		Description
Media		
	Media ID	The media ID for the volume. To search on all media, select <All> .
	Media host	The host name of the media server that produced the originals. To search all hosts, select All media hosts .
	Disk type	The disk type of the storage unit.
	Disk pool	The name of the disk pool. Not enabled if the disk type is BasicDisk.
	Media server	The name of the media server that produced the original images. To search all media servers, select All media hosts .
	Volume	The ID of the disk volume in the disk pool. Enabled if the disk type is not BasicDisk.
	Path	Searches for an image on a disk storage unit, if the path is entered. Or, searches all of the disk storage on the specified server, if All was selected. Enabled if the disk type is BasicDisk.
Date/time range		The range of dates and times that you want to search. The Global attributes property Policy update interval determines the default range.
Copies, policies, and clients		
	Copies	The copy that you want to search. Select either Primary or the copy number.
	Policy name	The policy under which the selected backups were performed. To search all policies, select All policies .
	Policy type	The purpose of the policy.
	Type of backup	The type of schedule that created the backup. To search all schedule types, select All backup types . Enabled if you select a specific Policy type .
	Client (host name)	The host name of the client that produced the backup. To search all hosts, select All clients .
Job priority		
	Override default job priority	The job priority for the catalog action (verify, duplicate, or import). To change the default, enable Override default priority. Then, select a value for the Job priority. If this option is not enabled, the job runs using the default priority as specified in the Default job priorities host property. Changes that you make affect the priority for the selected job only.

Table 24-1 Catalog search criteria (*continued*)

Property	Description
Job priority	The priority of the catalog job. Enabled if you override the default priority.

Search results properties

In addition to properties that you can select for the search, other properties are displayed for the images.

Table 24-2 Catalog search results properties

Property	Description
Copy DTE mode	Specifies whether the data is transferred over a secure channel when the current image copy is created.
Copy hierarchy DTE mode	Specifies whether the data is transferred over a secure channel when the current image copy and all its parent copies in the hierarchy are created.
Expiration date	The date that the image expires. This option is not available in the Administration Console.
Image DTE mode	Indicates the data-in-transit encryption (DTE) mode for the backup image.
Immutable	Indicates if the backup image is read-only and cannot be modified, corrupted, or encrypted.
Indelible	Indicates if the backup image is protected from being deleted before it expires.
Malware scan status	The scan status of the backup image.
Infection status	Displays the malware infected status of the backup images. The infection can be detected by malware scan or file hash search.
Mirror copy	Indicates if the image is a mirror replica or copy.

Table 24-2 Catalog search results properties (*continued*)

Property	Description
On hold	Indicates whether the image copy is on hold or not. Yes: The image has only one copy and a hold is set on the copy. No: No hold is set on the copy. A hold is set with the <code>nbholdutil</code> command.
Time	The time that the backup ran.
WORM unlock time	Indicates the time at which the image can be altered or deleted. Applies to the storage units that are WORM capable.

Verifying backup images

NetBackup can verify the contents of a backup by reading the volume and comparing its contents to what is recorded in the NetBackup catalog.

This operation does not compare the data on the volume to the contents of the client disk. However, the operation does read each block in the image to verify that the volume is readable. (However, data corruption within a block is possible.)

NetBackup verifies only one backup at a time and tries to minimize media mounts and positioning time.

To verify backup images

- 1 Open the NetBackup web UI.
- 2 On the left, click **Catalog**.
- 3 From the **Action** list, select **Verify**.
- 4 Select the search criteria to find the image you want to verify. Click **Search**.
Backups that have fragments on another volume are included, as they exist in part on the specified volume.
- 5 Select the image that you want to verify. Then click **Verify**.
- 6 Click the **Activity** tab to view the job results.

Promoting a copy to a primary copy


Each backup is assigned a primary copy. NetBackup uses the primary copy to satisfy restore requests. The first backup image that is created successfully by a NetBackup policy is the primary backup. If the primary copy is unavailable and a duplicate copy exists, select a copy of the backup and set it to be the primary copy.

NetBackup restores from the primary backup, and Vault duplicates from the primary backup. If your Vault profile performs duplication, you can designate one of the duplicates as the primary. In most circumstances, the copy remaining in the robot is the primary backup. When a primary backup expires, the next backup (if it exists) is promoted to primary automatically.

Use one of the following methods to promote a copy to a primary copy:

- Promote a backup copy to a primary copy
- See [the section called “Promote a backup copy to a primary copy”](#) on page 927.
- Promote a copy to a primary copy for many backups using the `bpchangeprimary` command
- See [the section called “Promoting a copy to a primary copy for many backups”](#) on page 928.

Figure 24-1 Primary copy status



Primary Copy status indicates that the image is now the primary copy

Promote a backup copy to a primary copy

To promote a backup copy to a primary copy

- 1 Open the NetBackup web UI.
- 2 On the left, click **Catalog**.
- 3 From the **Action** list, select **Duplicate**.
- 4 Select the search criteria to find the image you want to promote. Be sure that you indicate a copy in the **Copies** field and not **Primary copy**.
- 5 Click **Search**.

- 6 Select the image you want to promote. Then click **Set primary copy**.

After the image is promoted to the primary copy, the **Primary copy** column immediately reads **Yes**.

- 7 Click the **Activity** tab to view the job results.

Promoting a copy to a primary copy for many backups

More information on the `bpchangeprimary` is available in the [NetBackup Commands Reference Guide](#).

To promote a copy to a primary copy for many backups

- ◆ You can also promote a copy to be a primary copy for many backups using the `bpchangeprimary` command. For example, the following command promotes all copies on the media that belongs to the `b_pool` volume pool. The copies must have been created after August 8, 2022:

```
bpchangeprimary -pool b_pool -sd 08/01/2022
```

In the next example, the following command promotes copy 2 of all backups of `client_a`. The copies must have been created after January 1, 2022:

```
bpchangeprimary -copy 2 -cl client_a -sd 01/01/2022
```

Duplicating backup images

NetBackup does not verify in advance whether the storage units and the drives that are required for the duplicate operation are available for use. NetBackup verifies that the destination storage units exist. The storage units must be connected to the same media server.

[Table 24-3](#) lists the scenarios in which duplication is or is not possible:

Table 24-3 Backup duplication scenarios

Duplication possible	Duplication not possible
■ From one storage unit to another. ■ From one media density to another. ■ From one server to another. ■ From multiplex to nonmultiplex format. ■ From multiplex format and retain the multiplex format on the duplicate. The duplicate can contain all or any subset of the backups that were included in the original multiplexed group. The duplicate is created with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.)	■ While the backup is created (unless making multiple copies concurrently). ■ When the backup has expired. ■ By using NetBackup to schedule duplications automatically (unless you use a Vault policy to schedule duplication) ■ When it is a multiplexed duplicate of the following type: ■ FlashBackup ■ NDMP backup ■ Backups from disk type storage units ■ Backups to disk type storage units ■ Nonmultiplexed backups

An alternative to duplicating backups is to create up to four copies simultaneously at backup time. (This option is sometimes referred to as Inline Copy.) Another alternative is to use storage lifecycle policies.

See [“About writing multiple copies using a storage lifecycle policy”](#) on page 639.

To duplicate backup images

- 1 Open the NetBackup web UI.
- 2 On the left, click **Catalog**.
- 3 From the **Action** list, select **Duplicate**.
- 4 Select the search criteria to find the image you want to duplicate.
See [“Catalog utility search criteria and backup image details”](#) on page 923.
- 5 Select the images that you want to duplicate and click **Duplicate**.
If you duplicate a catalog backup, select all child jobs that were used to create the catalog backup. All jobs must be duplicated to duplicate the catalog backup.
- 6 Specify the number of copies you want to create. NetBackup can create up to 10 copies of unexpired backups.
If enough drives are available, the copies are created simultaneously. Otherwise, the system may require operator intervention if four copies are to be created using only two drives, for example.

- 7** The primary copy is the copy from which restores are done. Normally, the original backup is the primary copy.

If you want one of the duplicated copies to become the primary copy, select the copy number from the drop-down, otherwise select **Keep current primary copy**.

When the primary expires, a different copy automatically becomes primary. (The copy that is chosen is the one with the smallest copy number. If the primary is copy 1, copy 2 becomes primary when it expires. If the primary is copy 5, copy 1 becomes primary when it expires.)

- 8** Specify the storage unit where each copy is stored. If a storage unit has multiple drives, it can be used for both the source and destination.

All storage units must meet the criteria for creating multiple copies.

See [“About configuring multiple copies”](#) on page 747.

- 9** Specify the volume pool where each copy is stored.

The following volume pool selections are based on the policy type setting that was used for the query.

If the Policy type is set to All policy types (default).	Specifies that all volume pools are included in the drop-down list. Both catalog and non-catalog volume pools are included.

If the Policy type is set to NBU-Catalog .	Specifies that only catalog volume pools are included in the drop-down list.

If the Policy type is set to a policy type other than NBU-Catalog or All policy types .	Specifies that only non-catalog volume pools are included in the drop-down list.

NetBackup does not verify that the media ID selected for the duplicate copy is different from the media ID that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure that a different volume is used.

- 10** Select the retention level for the copy, or select **No change**.

The duplicate copy shares many attributes of the primary copy, including backup ID. Other attributes apply only to the primary. (For example, elapsed time.) NetBackup uses the primary copy to satisfy restore requests.

Consider the following items when selecting the retention level:

- If **No change** is selected for the retention period, the expiration date is the same for the duplicate and the source copies. You can use the `bpexptime` command to change the expiration date of the duplicate.
 - If a retention period is indicated, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2022 and its retention period is one week, the new copy's expiration date is November 21, 2022.
- 11** Specify whether the remaining copies should continue or fail if the specified copy fails.
- 12** Specify who should own the media onto which you duplicate images.

Select one of the following:

Any	Specifies that NetBackup chooses the media owner, either a media server or server group.
None	Specifies the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	Specifies that only those media servers in the group are allowed to write to the media on which backup images for this policy are written. All of the media server groups that are configured in your NetBackup environment appear in the drop-down list.

- 13** If the selection includes multiplexed backups and the backups are to remain multiplexed in the duplicate, select **Preserve multiplexing**. If you do not duplicate all the backups in a multiplexed group, the duplicate contains a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together during a single session.)

By default, duplication is done serially and attempts to minimize media mounts and positioning time. Only one backup is processed at a time. If **Preserved multiplexing** is enabled, NetBackup first duplicates all backups that cannot be multiplex duplicated before the multiplexed backups are duplicated.

The **Preserve multiplexing** setting does not apply when the destination is a disk storage unit. However, if the source is a tape and the destination is a disk storage unit, select **Preserve multiplexing** to ensure that the tape is read in one pass.

- 14 Click **Yes** to start duplicating.
- 15 Click the **Activity** tab, then select the duplication job to view the job results.
See [“Multiplexed duplication considerations”](#) on page 932.

Multiplexed duplication considerations

Consider the following items about multiplexed duplication.

Table 24-4 Multiplexed duplication considerations

Consideration	Description
Multiplex settings are ignored	When multiplexed backups are duplicated, the multiplex settings of the destination storage unit and the original schedule are ignored. However, if multiple multiplexed groups are duplicated, the grouping within each multiplexed group is maintained. This means that the duplicated groups have a multiplexing factor that is no greater than the factor that was used during the original backup.
Backups in a multiplexed group are duplicated and duplicated group is identical	When backups in a multiplexed group are duplicated to a storage unit, the duplicated group is identical as well. However, the storage unit must have the same characteristics as the unit where the backup was originally performed. The following items are exceptions: ■ If EOM (end of media) is encountered on either the source or the destination media. ■ If any of the fragments are zero length in the source backups, the fragments are removed during duplication. A fragment of zero length occurs if many multiplexed backups start at the same time.

Jobs that appear while making multiple copies

When multiple copies are made concurrently, a parent job appears, plus a job for each copy.

The parent job displays the overall status, whereas the copy jobs display the status of a single copy. Viewing the status of individual jobs lets you troubleshoot jobs individually. For example, if one copy fails but the other copy is successful, or if each copy fails for different reasons. If at least one copy is successful, the status of the parent job is successful. Use the **Parent Job ID** filter to display the parent Job ID. Use the **Copy number** filter to display the copy number for a particular copy.

Expiring backup images

To expire a backup image means to force the retention period to expire, or information about the backup is deleted. When the retention period expires, NetBackup deletes information about the backup. The files in the backups are unavailable for restores without first re-importing.

To expire a backup image

- 1 Open the NetBackup web UI.
- 2 On the left, click **Catalog**.
- 3 Select the search criteria to find the image you want to duplicate.
See [“Catalog utility search criteria and backup image details”](#) on page 923.
- 4 Select the image you want to expire and click **Expire > Expire**.

About Image Dependency Expiration Cleanup

Image Dependency Expiration Cleanup checks whether any other backup images depend on an expiring backup image. It runs as part of an image clean-up job. The feature prevents a backup image from expiring if it has any subsequent dependent backup images.

For example, when it expires a new FULL backup image, NetBackup checks whether there are any backup images that depend on the FULL image. (These dependent images include differential incremental (INCR) or cumulative incremental (CINCR) images. If the expiring FULL image has dependencies, it is put on hold and blocked from expiring until all dependencies expire or are removed. Similarly, an expiring incremental backup image with any backup image dependencies is put on hold until all the dependencies expire or are removed. This process can lead to holding a chain of dependent images. Consider an example where a chain of 10 expired images is held for a dependent copy. In that case it takes 10 image cleanup jobs to process and remove all of the images in the chain.

Explicit expiration (either with `bpexpdate` or the **Expire** action in the web UI) does not perform this check and therefore does not block an image from expiring.

Image Dependency Expiration Cleanup applies only to new backup images created with one of these supported policies:

- Cloud
- CloudStorage
- Kubernetes
- MS-Windows

- NDMP
- Hyper-V
- HyperScale
- Hypervisor
- NAS-Data-Protection
- Standard
- VMware
- Nutanix-AHV
- MongoDB Ops Manager

Note: When you configure schedules, keep in mind that "forever" incremental backups might lead to holding all the images perpetually because each image depends on other image. So, it is recommended that you configure schedules in combination with FULL schedules.

Checking whether an image is "held for dependent copy"

Images that are put on hold due to Image Dependency Expiration Cleanup are marked as **Held for dependent copy** and reported at the copy level.

To see whether an image copy is held, enter the following command:

```
bpimagelist -backupid <backup_ID> -[-l|-L|-json]
```

Review the output for the following statement:

```
Held For Dependent Copy: yes
```

You can also check the NetBackup web UI (**Catalog > <backup_image>**) to see whether an image copy is held.

Enabling and disabling the feature

Note: Image Dependency Expiration Cleanup is enabled by default. When it is enabled, image cleanup jobs can run up to 2.5 times slower than when the feature is disabled.

To see whether Image Dependency Expiration Cleanup is enabled, enter the following command:

```
bpconfig -U
```

Review the output for the following statement:

Image Expiry Dependency Check: (enabled)

To disable Image Dependency Expiration Cleanup, enter the following command:

```
bpconfig -image_expiry_dependency_check 0
```

To turn on the feature, enter the following command:

```
bpconfig -image_expiry_dependency_check 1
```

You can also use a REST API to turn off or turn on the feature. See the API documentation for more information.

About importing backup images

NetBackup can import the backups that have expired or the backups from another NetBackup server.

During an import operation, NetBackup recreates NetBackup catalog entries for the backups on the imported volume. The import capability is useful for moving volumes from one site to another and for recreating NetBackup catalog entries.

An image is imported in the following two phases:

Table 24-5 Phases to import an image

Phase	Description
Phase I: Initiate Import	NetBackup creates a list of expired catalog entries for the backups on the imported volume. No actual import occurs in Phase I. See “Importing backup images, Phase I” on page 936.
Phase II: Import	Images are selected for importing from the list of expired images that was created in Phase I. See “Importing backup images, Phase II” on page 937.

About importing expired images

The expiration date for the imported items is the current date plus the retention period. For example, if a backup is imported on November 14, 2021, and its retention period is one week, the new expiration date is November 21, 2021.

Consider the following items when importing backup images:

- NetBackup can import the disk images that NetBackup version 6.0 (or later) writes.
- You cannot import a backup if an unexpired copy of it already exists on the server.

- NetBackup does not direct backups to imported volumes.
- If you import a catalog backup, import all the child jobs that were used to create the catalog backup. All jobs must be imported to import the catalog backup.
- To import a volume with the same media ID as an existing volume on a server, use the following example where you want to import a volume with media ID A00001. (A volume with media ID A00001 already exists on the server.)
 - Duplicate the existing volume on the server to another media ID (for example, B00001).
 - Remove information about media ID A00001 from the NetBackup catalog by running the following command:
 On Windows:

```
install_path\NetBackup\bin\admincmd\bpxpdate
-d 0 -m mediaID
```

 On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpxpdate -d 0 -m
media_ID
```
 - Delete media ID A00001 from Media Manager on the server.
 - Add the other A00001 to Media Manager on the server.

To avoid this problem in the future, use unique prefix characters for media IDs on all servers.

See [“Expiring backup images”](#) on page 933.

Importing backup images, Phase I

Phase I of the import process creates a list of images from which to select to import in Phase II. No import occurs in Phase I.

Note the following about importing backup images:

- If tape is used, each tape must be mounted and read. It may take some time to read the catalog and build the list of images.
- The backup is not imported if it begins on a media ID that the initiating backup procedure did not process.
- The backup is incomplete if it ends on a media ID that the initiating backup procedure did not process.
- To import a catalog backup, import all of the child jobs that were used to create the catalog backup.

To perform Phase I: initialize import of backup images

- 1 To import the images from tape, make the media accessible to the media server so the images can be imported.
- 2 Open the NetBackup web UI.
- 3 On the left, click **Catalog**.
- 4 On the **Actions** menu, select **Phase I import**.
- 5 For the **Media server**, specify the name of the host that contains the volume to import. This media server becomes the media owner.
- 6 Indicate the location of the image. For the **Image type**, select whether the images to be imported are located on tape or on disk.

The following table shows the actions to take depending on the location of the image.

If images are on tape	In the Media ID field, enter the Media ID of the volume that contains the backups to import.
If images are on disk	In the Disk type field, select the type of the disk storage unit on which to search for backup images. The disk types depend on which NetBackup options are licensed. If the disk type references a disk pool, enter or select the disk pool and the disk volume ID. For a BasicDisk type, enter or browse to the path to the images in the field provided. For other disk types, select <All> or the specific volume.

- 7 Click **Import** to begin reading the catalog information from the source volume.
- 8 Click on the **Activity** tab to watch as NetBackup looks at each image on the tape. NetBackup determines whether or not each image has expired and can be imported. The job also displays in the Activity monitor as an Image import type. Select the import job log to view the job results.

Importing backup images, Phase II

To import the backups, first run the Initiate Import operation (Import Phase I). The first phase reads the catalog to determine all of the media that contain the catalog backup images. After Phase I, start the Import operation (Phase II). If Phase II is run before Phase I, the import fails with a message. For example, Unexpected EOF or Import of backup ID failed, fragments are not consecutive.

To import backup images, Phase II

- 1 Open the NetBackup web UI.
- 2 On the left, click **Catalog**.
- 3 On the **Actions** menu, select **Phase II import**.
- 4 Set up the search criteria to find images available to import. Be sure to select a date range that includes the images you want to import. Click **Search**.
- 5 Select the images that you want to import. Click **Import** to import the selected images.
- 6 Select whether you'd like to log the names of all of the files that are found in the imported images. Click **OK**.
- 7 Click the **Activity** tab to view the progress of Import phase II.

Configuring immutability and indelibility of data in NetBackup

This chapter includes the following topics:

- [About immutable and indelible data](#)
- [Workflow to configure immutable and indelible data](#)
- [Deleting an immutable image from storage using the bpexpdate command](#)
- [Removing an immutable image from the catalog using the bpexpdate command](#)

About immutable and indelible data

NetBackup protects your data from being encrypted, modified, and deleted using WORM properties.

WORM is the acronym for Write Once Read Many.

WORM properties provide two additional levels of security for backup images:

- **Immutability** - this protection ensures that the backup image is read-only and cannot be modified, corrupted, or encrypted after backup.
- **Indelibility** - this property protects the backup image from being deleted before it expires. The data is protected from malicious deletion.

Configuring these WORM properties protects your data from certain malware attacks to some extent, for example ransomware.

NetBackup provides the ability to write backups to WORM storage devices so their data cannot be corrupted. Additionally, it lets you take advantage of advanced

options available from your storage vendors to ensure backups are retained unaltered on storage platforms to meet regulatory and compliance requirements.

All NetBackup image copies have an Expiration Time. This time is calculated by using the configured retention level in the schedule and the start time of the backup job.

When a NetBackup image is written to a WORM-enabled storage unit, the data cannot be altered or deleted until the WORM Unlock Time for that image has elapsed. Unlike the Copy Expiration time that is calculated from the start time of the backup job, the WORM Unlock Time is associated with the WORM storage. The WORM Unlock Time value is calculated using the configured retention level and the write completion timestamp for the backup image onto WORM storage.

When you use `bpimagelist` to view an image that is written to WORM storage, the timestamp that is associated with the Copy Expiration time precedes the WORM Unlock Time for that copy of the backup image. For longer-running backups or duplication jobs, the difference is greater between Copy Expiration Time and WORM Unlock Time.

As part of normal operations, copies of backup images on WORM storage are not removed from the catalog and storage until both Copy Expiration Time and Worm Unlock Time timestamps have elapsed. The WORM Unlock Time of a copy that is written to WORM storage can only be extended and cannot be shortened. To extend the expiration date, use the `bpexpdate -extend_worm_locks` command.

In special circumstances, the `bpexpdate -try_expire_worm_copy` option can be used to force an attempted removal of a WORM indelible image from the NetBackup catalog. This option is only recommended to be used after removing WORM locks directly on the storage device. Only use this option with assistance from Cohesity technical support.

When duplicating an image onto WORM storage, the WORM Unlock Time can be configured to match the Copy Expiration Time by running the `bpduplicate` command using the `-worm_unlock_match_expiration` option that was introduced in NetBackup 10.1.

If older backup images are duplicated to WORM storage without using this command option, the WORM Unlock Time for the duplicated copy is calculated using the configured retention level, and the timestamp when the duplication job was complete.

The `bpduplicate -worm_unlock_match_expiration` command option is not used for SLP driven duplications. For SLP driven duplications, the retention period is applied from the end of the duplication job to calculate WORM Unlock Time of the new copy. The Copy Expiration Time for the new copy is calculated from the retention period that is applied to the backup time (for copy 1).

For A.I.R. jobs, the retention period is applied from the end of the import job to calculate the WORM Unlock Time of the imported copy. The Copy Expiration Time is calculated as the retention period that is applied from the beginning of the import job.

For more information about the `bpduplicate` command and the `bpexpdate` command, see the [NetBackup Commands Reference Guide](#).

Note: When you use the `bpduplicate -worm_unlock_match_expiration` and `bpexpdate -extend_worm_locks` command options, they rely on the accuracy of the NetBackup primary server clock. That is because the WORM Unlock Time mirrors the Image Expiration timestamp for that copy.

For more information about how to base the WORM Unlock Time on the original backup time, see the following knowledge base article:

[Images duplicated to WORM storage have unlock time calculated from duplication date not backup date](#)

Workflow to configure immutable and indelible data

Carry out the following steps in the given order to protect your data by configuring immutability and indelibility.

Table 25-1 Workflow to configure immutable and indelible data

Step	Description
1	Configure the following WORM settings on the storage server. The storage administrator configures these settings outside of NetBackup. ■ WORM capable - If the storage unit and the associated disk pool are enabled to use the WORM property at the time of backup image creation, the backup images are set to be immutable and indelible. ■ Lock Minimum Duration - Specifies the minimum allowed duration for which the data for a backup image is indelible. The storage administrator sets this duration on the Logical Storage Unit (LSU) or the Domain Volume (DV), which NetBackup discovers. ■ Lock Maximum Duration - Specifies the maximum allowed duration for which the data for a backup image is indelible. The storage administrator sets this duration on the Logical Storage Unit (LSU) or the Domain Volume, which NetBackup discovers. Refer to the OST vendor plug-in documentation.

Table 25-1 Workflow to configure immutable and indelible data (*continued*)

Step	Description
2	Configure a disk pool using WORM-capable volumes. See “About configuring disk pool storage” on page 406.
3	Configure a storage unit with the Use WORM option enabled. See “Use WORM setting” on page 574.
4	Configure a backup policy using the WORM-enabled storage unit. See “Creating a backup policy” on page 657.

Note: In case of storage changes or third-party OST vendor software upgrades, you need to manually update the storage servers and the disk pools. See the 'Completing your system update after an upgrade' section from the [NetBackup Upgrade Guide](#).

Deleting an immutable image from storage using the `bpexpdate` command

Deletion of an immutable image can only happen when storage is used that allows for lock deletion. The lock deletion can be done using the Enterprise mode on Flex Appliance, Flex Scale Appliance, Access Appliance, or a third-party storage device that supports lock deletion. When an immutable image is deleted, the storage that you use is responsible for the lock deletion and NetBackup is responsible for the image deletion.

When you use Flex Appliance, Flex Scale Appliance, or Access Appliance, you must use the command line or an SSH session to remove the lock on the image. If you use a third-party storage device, refer to that vendor’s documentation for steps on removing locked images.

To delete the immutable image on the appliance

- 1 Verify that the appliance is in Enterprise mode.
- 2 From the NetBackup Command Line, use `bpimagelist` command to find the image ID.

This procedure uses the following example image ID:

```
Backup ID: server123.veritas.com_1234567890
```

- 3 Delete the image lock on storage using the command line option or the SSH session option.
 - For Flex Appliance: You must use the default `msdpadm` user to run the following options.
 - For Flex Scale Appliance and Access Appliance: You must use an appliance user with the appliance administrator role.

Command line option:

- Open the `/usr/openv/pdde/pdcr/bin/` directory.
- Use the following command to query and modify the catalog database for the given backup ID (Example: `server123.veritas.com_1234567890`). The `-worm disable` option disables the retention lock for an image using the backup ID.

```
sudo -u msdpvc /usr/openv/pdde/pdcr/bin/catdbutil -worm
disable -backupid
```

SSH session option:

- Open an SSH session to the WORM storage server instance.
- Use the `retention policy disable` command to query and modify the catalog database for the given policy. The `policydisable` arguments disable the retention lock for an image using the policy ID used for the image retention that has a retention lock.

For more information about the command options in this step, see the [NetBackup Deduplication Guide](#).

- 4 Add the image ID to `bpexpdate` with the `-try_expire_worm_copy` option.

```
bpexpdate -d 0 backupid server123.veritas.com_1234567890
-try_expire_worm_copy -copy 1
```

- 5 Use `y` or `n` to confirm deletion.

If the storage lock is not removed, NetBackup returns an error indicating that there is a WORM lock error.

See [“Removing an immutable image from the catalog using the bpexpdate command”](#) on page 944.

See [“About immutable and indelible data”](#) on page 939.

Removing an immutable image from the catalog using the `bpexpdate` command

You can remove an immutable image from the NetBackup catalog and have that image remain on storage.

To remove an immutable image from the catalog

- 1 Open the NetBackup Command Line Interface (CLI).
- 2 Delete the image from the catalog using the `bpexpdate` command with the `-try_expire_worm_copy` and the `-nodelete` options.

```
bpexpdate -d 0 -backupid server123.veritas.com_1234567890  
-copy 1 -try_expire_worm_copy -nodelete
```

Using the `-try_expire-worm_copy` and `-nodelete` options together removes the image from the catalog only and does not affect storage.

- 3 Use `y` or `n` to confirm deletion.

See [“Deleting an immutable image from storage using the `bpexpdate` command”](#) on page 942.

See [“About immutable and indelible data”](#) on page 939.

Deployment Management

- [Chapter 26. Deployment Management](#)

Deployment Management

This chapter includes the following topics:

- [About the deployment policies utility](#)
- [Deployment policy management](#)
- [Copying a deployment policy to create a new deployment policy](#)
- [Attributes tab in Deployment management](#)
- [Schedules tab in Deployment management](#)
- [Security options tab in Deployment management](#)
- [Manually deploy a deployment policy](#)
- [Perform client initiated upgrade with VxUpdate](#)
- [Deployment job status](#)

About the deployment policies utility

Deployment policies are the main component of VxUpdate that serves as a client or host upgrade tool. The deployment policy lets you configure and run deployment activities on a schedule or enable the host owners to upgrade at their convenience. You can schedule precheck, staging, and installation tasks as separate activities with different schedules, each with their own specific deployment windows.

For more information regarding VxUpdate, see the *About VxUpdate* section within the [NetBackup Upgrade Guide](#).

Deployment policies are located in the NetBackup web UI under **Hosts > Deployment Management**. These policies provide the instructions that NetBackup follows to upgrade clients or hosts. Use this utility to provide the following instructions for a client or host upgrade:

What type of client or host to upgrade	See “Attributes tab in Deployment management” on page 952.
Which clients or hosts to upgrade	
When to perform VxUpdate	See “Schedules tab in Deployment management” on page 953.
The security options to use for the clients or hosts	

Deployment policy management

Use the procedures that are shown to create, modify, and delete your deployment policies.

Creating a deployment policy

Note: You must add packages to the VxUpdate repository before you can create a working deployment policy. You can create deployment policies without packages in the repository, but those policies fail to run successfully.

For more information regarding adding packages, see the *Repository Management* section within the [NetBackup Upgrade Guide](#).

- 1 In the NetBackup Administration Console, in the left pane, select **Deployment Management > Deployment Policies**.
- 2 From the **Actions** menu, select **New Deployment Policy**.
- 3 Enter a unique name for the new policy in the **Add a New Deployment Policy** dialog box.
- 4 Click **OK**.
- 5 Specify the information that is shown on the **Attributes** tab in the **Change Deployment Policy** window:
 - **Package:** Select the package that you want deployed from the drop-down menu. Only packages that are in the **Ready** state are available for use.

Note: Specifying a package that supports external certificate authority certificates presents you with an additional tab titled **Security**. That tab is covered later in this procedure.

- **Media server:** Specify the media server from drop-down. The media server that is specified is used to connect and transfer files to the NetBackup hosts that are included in the policy. The media server also caches the files from the NetBackup repository. Since the repository resides on the primary server, the primary server is the default value for the media server field.
- **Java GUI and JRE:** Specify if you want the Java GUI and the JRE upgraded on the target systems. The three options include:
 - **INCLUDE:** Install or upgrade the Java GUI and JRE components on the specified computers.
 - **EXCLUDE:** Exclude the Java GUI and JRE components from the specified computer. Any preexisting NetBackup Java GUI and JRE packages are removed.
 - **MATCH:** Preserve the current state of the Java GUI and JRE components. The components are upgraded if they are present on the pre-upgraded system. The components are not installed if they are not present on the pre-upgraded system.
- **NetBackup DirectIO:** Specify if you want the NetBackup DirectIO package upgraded on the target systems. The three options include:
 - **INCLUDE:** Install or upgrade the NetBackup DirectIO package on the specified computers.
 - **EXCLUDE:** Exclude the NetBackup DirectIO package from the specified computer. Any preexisting NetBackup DirectIO package binaries are removed.
 - **MATCH:** Preserve the current state of the NetBackup DirectIO package. The plug-in is upgraded if it is present on the pre-upgraded system. The plug-in is not installed if they are not present on the pre-upgraded system.
- (Conditional): Select the **Limit simultaneous jobs** option and specify a value for **jobs** to limit the total number of concurrent jobs that can run at a time. The minimum value is 1 and the maximum value is 999. If the check box is selected, the default value is 3. If you do not select the check box, no limit is enforced for the simultaneous upgrade jobs. You can set unlimited simultaneous upgrade jobs through command line interface by setting the value as 0.
- **Select hosts:** Select hosts from the **Available hosts** list and select **Add** to add hosts to the deployment policy. The list is generated from hosts in the host database and backup policies. Once you select **Add**, the hosts are shown under **Selected hosts**.

- 6 Select the **Schedules** tab in the **Change Deployment Policy** window.

You can see a summary of all schedules within that policy.

- 7 Select **New**.

- 8 Specify the information that is shown in the **Add Deployment Schedule** window.

- **Name:** Enter a name for the new schedule.
- **Type:** Specify the type of schedule you want created.

Schedule types:

- **Precheck**
Performs the various precheck operations, including confirming there is sufficient space on the client for the update. The precheck schedule type does not exist for EEB packages.
- **Stage**
Moves the update package to the client, but does not install it. Also performs the precheck operation.
- **Install**
Installs the specified package. Also performs the precheck and the stage package operations. If you already performed the stage package operation, the install schedule does not move the package again.

Note: Please be aware that adding multiple different schedule types to the same deployment schedule window has unpredictable results. VxUpdate has no defined behavior to determine which schedule type runs first. If a single deployment schedule window has precheck, stage, and install jobs, there is no way to specify the order in which they run. The precheck or the stage schedules can fail, but the install completes successfully. If you plan to use precheck, stage, and install schedules, it is recommended that you create separate schedules and separate windows for each.

- **Starts:** Specify the date and time you want the policy to start in the text field or with the date and the time spinner. You can also click the calendar icon and specify a date and time in the resulting window. You can select a schedule by clicking and dragging over the three-month calendar that is provided at the bottom of the window.
- **Ends:** Specify the date and time you want the policy to end as you specified the start time.

- **Duration:** Optionally, you can specify a duration in days, hours, minutes, and seconds instead of an end time for the policy. The minimum value is 5 minutes and the maximum is 99 days.
- Select **Add/OK** and the schedule is created. Select **OK** to save and create your policy.

- 9 A **Security** tab appears when you select a deployment package that contains support for external certificate authorities.

By default, the **Use existing certificates when possible** option is selected. This option instructs NetBackup to use the existing NetBackup CA or external CA certificates, if available.

Note: If you specify this option and certificates are not available, your upgrade fails.

Deselecting the **Use existing certificates when possible** option lets you specify the location for external certificate authority information for both UNIX and Linux computers and Windows computers.

Deselecting this option does not allow the user to change the security configuration settings during the upgrade.

- 10 Windows clients have **Use Windows certificate store** selected by default.

You must enter the certificate location as *Certificate Store Name\Issuer Distinguished Name\Subject Distinguished Name*.

Note: You can use the `$hostname` variable for any of the names in the certificate store specification. The `$hostname` variable evaluates at run time to the name of the local host. This option provides flexibility when you push NetBackup software to a large number of clients.

Alternatively, you can specify a comma-separated list of Windows certificate locations. For example, you can specify:

```
MyCertStore\IssuerName1\SubjectName,  
MyCertStore\IssuerName2\SubjectName2,  
MyCertStore4\IssuerName1\SubjectName5
```

Then select the Certificate Revocation List (CRL) option from the radio buttons shown:

- **Do not use a CRL.** No additional information is required.
- **Use the CRL defined in the certificate.** No additional information is required.

- **Use the CRL at the following path:** You are prompted to provide a path to the CRL.
- 11 ■ **Certificate file:** This field requires you to provide the path to the certificate file and the certificate file name.
- **Trust store location:** This field requires you to provide the path to the trust store and the trust store file name.
- **Private key path:** This field requires you to provide the path to the private key file and the private key file name.
- **Passphrase file:** This field requires you to provide the path of the passphrase file and the passphrase file name. This field is optional.
- Then specify the correct CRL option for your environment:
 - **Do not use a CRL.** No additional information is required.
 - **Use the CRL defined in the certificate.** No additional information is required.
 - **Use the CRL at the following path:** You are prompted to provide a path to the CRL.

To change a deployment policy

- 1 Right click on the deployment policy and select **Change**.
- 2 Navigate through the deployment policy tabs and make any necessary changes to the policy.
- 3 Select **OK** and the policy is updated.

Deleting a deployment policy

- 1 Right click on the deployment policy and select **Delete**.
- 2 Select **OK**.
- 3 Confirm the deletion of the policy.

Copying a deployment policy to create a new deployment policy

Use the **Copy policy** option to save time creating policies. This option is especially useful for the policies that contain many of the same policy attributes, schedules, or hosts selections.

To copy a deployment policy

- 1 Open the NetBackup web UI.
- 2 On the left, select **Hosts > Deployment management**.
- 3 Select the policy to copy.
- 4 Select **Copy policy**.
- 5 Enter the name for the new policy.
- 6 Select **Copy**. The only difference between the new policy and the copied policy is the name.

Make any changes that you want to the new policy. Then select **Copy**.

Attributes tab in Deployment management

Use the policy **Attributes** tab to configure deployment management settings when you add a new deployment policy or change an existing deployment policy.

Setting	Description
Package	Select the package that you want to deploy. Note: You must add packages to the VxUpdate repository before you can create a working deployment policy. You can create deployment policies without packages in the repository, but those policies fail to run successfully. For more information regarding adding packages, see the <i>Repository Management</i> section within the NetBackup Upgrade Guide.
Media server	Specify the media server. This media server is used to connect and transfer files to the NetBackup hosts that are included in the policy. The media server must be version NetBackup 8.1.2 or later. Since the repository resides on the primary server, the primary server is the default value for the media server field.

Setting	Description
Limit simultaneous jobs	Select the Limit simultaneous jobs option and specify a value for jobs to limit the total number of concurrent jobs that can run at one time. The default value is 3. The minimum value is 1 and the maximum value is 999. If you want to set unlimited simultaneous upgrade jobs, you must specify a value which is equivalent or higher than the count of the number of hosts that are selected for upgrade. For example, if you have selected 50 hosts, ensure that the Limit simultaneous jobs value is set to 50 or more but lower than the maximum value which is 999.
Java GUI and JRE	Specify if you want the NetBackup Administration Console and the JRE upgraded on the target systems. The three options include: ■ Match: Preserve the current state of the NetBackup Administration Console and JRE components. The components are upgraded if they are present on the pre-upgraded system. The components are not installed if they are not present on the pre-upgraded system. ■ Include: Install or upgrade the NetBackup Administration Console and JRE components on the specified computers. ■ Exclude: Exclude the NetBackup Administration Console and JRE components from the specified computer. Any preexisting NetBackup Administration Console and JRE packages are removed.

Schedules tab in Deployment management

Use the **Schedules** tab in Deployment management for the following tasks:

- To view a summary of all schedules within that policy.
- To create a new schedule.
- To edit or delete an existing schedule.

The schedules that are defined on the **Schedules** tab determine when VxUpdate occur for the selected deployment policy. The calendar displays a summary of all the schedules.

The **Schedules** tab contains both schedule information and other configuration options, beyond when the job is to run.

Setting	Description
Name	Enter a name for the new schedule.

Setting	Description
Operation	Specify the type of operation that you want to associate with the schedule. Precheck - Performs the various precheck operations, including confirming there is sufficient space on the client for the update. The precheck schedule type does not exist for EEB packages. Stage - Moves the update package to the client, but does not install it. This operation also performs the precheck operation. Install - Installs the specified package. This operation also performs the precheck and the stage package operations. If you already performed the stage package operation, the install schedule does not move the package again. Note: Be aware that adding multiple different schedule types to the same deployment schedule window has unpredictable results. VxUpdate has no defined behavior to determine which schedule type runs first. If a single deployment schedule window has precheck, stage, and install jobs, there is no way to specify the order in which they run. The precheck or the stage schedules can fail, but the install completes successfully. If you plan to use precheck, stage, and install schedules, it is recommended that you create separate schedules and separate windows for each.
Start date	Specify the date and time you want the policy to start in the text field or with the date and the time spinner. You can also click the calendar icon and specify a date and time in the resulting window. You can select a schedule by clicking and dragging over the three-month calendar that is provided at the bottom of the window.
End date	Specify the date and time you want the policy to end as you specified the start time.

Security options tab in Deployment management

Use the policy **Security options** tab to configure the settings for external security certificates. These settings are only available if a selected host is configured to use external certificates (certificates that are signed by a CA other than the NetBackup CA).

Attribute	Description
Use existing certificates when possible	This option instructs NetBackup to use the existing NetBackup CA or external CA certificates, if available. By default, the Use existing certificates when possible option is selected. Deselecting the Use existing certificates when possible option lets you specify the location for external certificate authority information for both UNIX and Linux computers and Windows computers. Note: If you specify this option and certificates are not available, your upgrade fails.
From Windows certificate store (Only for Windows)	Specifies that the certificate from the Windows certificate store is used. The certificate is searched using the following details that are provided with the Certificate location: Store name, Issuer name, Subject name.
Certificate file	Specifies the path to the external certificate of the host.
Trust store location	Specifies the path to the pem bundle of the Certificate Authorities.
Private key file	Specifies the path to the private key for the external certificate of the host.
Passphrase file	Specifies the path to the text file where the passphrase for the external certificate's private key is stored.
CRL check level	Specifies the revocation check level for the external certificate. It also lets you disable the revocation check for the external certificates. Based on the check level, the status of the certificate is validated against the Certificate Revocation List (CRL) during host communication. You can choose to use the CRLs from the directory that is specified in the NetBackup configuration file or the CRL Distribution Point (CDP).
From certificate file path (for file-based certificates) (Only for Windows)	Specifies a list of comma-separated clauses where each clause element contains a query. The clause is of the form <i><Store name>\<Issuer Name>\<Subject Name></i>. <i>\$hostname</i> is a keyword that is replaced with the fully qualified domain name of the host. For certificate selection from the Windows certificate store, NetBackup can pick a certificate from any of the Local Machine certificate stores on a Windows host. ■ Store name – The certificate store where the certificate is present ■ Issuer name (optional) – The certificate issuer name ■ Subject name – The subject name of the certificate If the issuer name is not specified, the certificate is searched based on the subject name.

Manually deploy a deployment policy

You can manually initiate a deployment policy based on an existing policy. Manually initiate deployment policies when you are logged into the server locally and need to force an immediate update. Or you can initiate an immediate upgrade for emergency binaries.

Use the **Deploy now** option to initiate a deployment job manually.

To manually deploy a deployment policy

- 1 Open the NetBackup web UI.
- 2 Go to **Hosts > Deployment management**. Then click the **Deployment policies** tab.
- 3 Select the policy that you want to start, and select **Deploy now**.
- 4 Select the operation that you want to run and the hosts that you want to upgrade.

If you do not select any hosts, NetBackup upgrades all hosts.
- 5 Click **Deploy now** to start the manual deployment job.

Perform client initiated upgrade with VxUpdate

Manually initiate deployment jobs when you are logged into the server locally and want to force an immediate update. You can also use a deployment job to initiate an immediate upgrade for emergency binaries.

Among the reasons for a client initiated upgrade using VxUpdate is mission critical systems with specific maintenance windows. One example of these systems is database servers with limited available down time.

To perform client initiated upgrade manually from the administration console

- 1 In the NetBackup Administration Console, navigate to **NetBackup Management > Host Properties > Clients**
- 2 In the right pane, right-click the client that you want to upgrade. Select **Upgrade Host**.
- 3 In the **Upgrade Host** dialog box, update the following fields:
 - **Package:** Select the package that you want to deploy from the drop-down menu.
 - **Type:** Select the deployment type from the drop-down menu.
 - Precheck

Note: The precheck schedule type does not exist for EEB packages.

- Stage
- Install
- **Media server:** Specify the media server from drop-down.
The media server must be version NetBackup 8.1.2 or later. Since the repository resides on the primary server, the primary server is the default value for the media server field.
- **Selected hosts:** Displays the list of selected hosts.

You can also perform client initiated upgrade using the command line option. For more information, refer *Perform client initiated upgrade with VxUpdate* section within the NetBackup Upgrade Guide.

Deployment job status

Monitor and review deployment job status in the Activity monitor. The **Deployment** job type is the new type for VxUpdate policies. Deployment policy parent jobs that exit with a status code 0 (zero) indicate that all the child jobs successfully completed. Parent jobs that finish with a status code 1 indicate that one or more of the child jobs succeeded, but at least one failed. Any other status code indicates failure. Review the status of the child jobs to determine why they failed. Otherwise, there are no differences between deployment jobs and other NetBackup jobs.

Your deployment job may receive a status code 224. This error indicates that the client's hardware and operating system are specified incorrectly. You can correct this error by modifying the deployment policy with the `bpplclients` command found in:

Linux: `/usr/opensv/netbackup/bin/admincmd`

Window: `install_path\netbackup\bin\admincmd`.

Use the syntax shown:

```
bpplclients deployment_policy_name -modify client_to_update -hardware  
new_hardware_value -os new_os_value
```

Deployment policies use a simplified naming scheme for operating system and hardware values. Use the values as shown for the `bpplclients` command:

Table 26-1 Deployment policy operating system and hardware

Operating system	Hardware
debian	x64
redhat	x64
suse	x64
redhat	ppc64le
suse	ppc64le
redhat	zseries
suse	zseries
aix	rs6000
solaris	sparc
solaris	x64
windows	x64

Security certificates are not deployed as part of the VxUpdate upgrade if the **Security level for certificate deployment** is set to **Very High**. This setting is located in the **Global security** settings.

If you cannot communicate with your clients after you use VxUpdate to upgrade your clients, please ensure that the proper security certificates were issued during upgrade. You may need to manually deploy the certificates. Refer to the following article that is shown for additional details:

https://www.veritas.com/content/support/en_US/article.100039650

Configuring replication

- [Chapter 27. About NetBackup replication](#)

About NetBackup replication

This chapter includes the following topics:

- [About NetBackup replication](#)
- [About NetBackup Auto Image Replication](#)
- [About NetBackup Replication Director](#)

About NetBackup replication

NetBackup offers two forms of replication:

Backups	Auto Image Replication Use this type of replication to replicate backups from one NetBackup domain to the NetBackup media server in another domain. See “About NetBackup Auto Image Replication” on page 961.
Snapshots	NetBackup Replication Director This type of replication makes use of NetBackup OpenStorage to replicate snapshots on primary storage to the disk arrays of OpenStorage partners. See “About NetBackup Replication Director” on page 1002. For more information, see the NetBackup Replication Director Solutions Guide.

About NetBackup Auto Image Replication

The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as Auto Image Replication.

The ability to replicate backups to storage in other NetBackup domains, often across various geographical sites, helps facilitate the following disaster recovery needs:

- **One-to-one model**
A single production data center can back up to a disaster recovery site.
- **One-to-many model**
A single production data center can back up to multiple disaster recovery sites. See [“One-to-many Auto Image Replication model”](#) on page 963.
- **Many-to-one model**
Remote offices in multiple domains can back up to a storage device in a single domain.
- **Many-to-many model**
Remote data centers in multiple domains can back up multiple disaster recovery sites.

NetBackup supports the following storage types for Auto Image Replication:

Table 27-1 NetBackup Auto Image Replication storage types

Storage type	Link to more information
Media Server Deduplication Pool	See the NetBackup Deduplication Guide .
An OpenStorage disk appliance	If your storage vendor's product supports replication, you can automatically replicate backup images to a similar device in a different primary server domain. See the NetBackup OpenStorage Solutions Guide for Disk.

Notes about Auto Image Replication

- Auto Image Replication does not support synthetic backups or optimized synthetic backups.
- Auto Image Replication does not support spanning volumes in a disk pool. NetBackup fails backup jobs to the disk pools that span volumes if the backup job is in a storage lifecycle policy that also contains a replication operation.

- Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group.
- The ability to perform Auto Image Replication between different versions of NetBackup does not overrule the basic image compatibility rules. For example, a database backup that was taken in one NetBackup domain can be replicated to a NetBackup domain of an earlier version. However, the older server may not be able to successfully restore from the newer image.
 For information about version compatibility and interoperability, see the NetBackup Enterprise Server and Server Software Compatibility List at the following URL:
<http://www.netbackup.com/compatibility>
- Synchronize the clocks of the primary servers in the source and the target domains so that the primary server in the target domain can import the images as soon as they are ready. The primary server in the target domain cannot import an image until the image creation time is reached. Time zone differences are not a factor because the images use Coordinated Universal Time (UTC).

Process Overview

[Table 27-2](#) is an overview of the process, generally describing the events in the originating and target domains.

NetBackup uses storage lifecycle policies in the source domain and the target domain to manage the Auto Image Replication operations.

See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 977.

Table 27-2 Auto Image Replication process overview

Event	Domain in which event occurs	Event description
1	Originating primary server (Domain 1)	Clients are backed up according to a backup policy that indicates a storage lifecycle policy as the Policy storage selection. The SLP must include at least one Replication operation to similar storage in the target domain.
2	Target primary server (Domain 2)	The storage server in the target domain recognizes that a replication event has occurred. It notifies the NetBackup primary server in the target domain.
3	Target primary server (Domain 2)	NetBackup imports the image immediately, based on an SLP that contains an import operation. NetBackup can import the image quickly because the metadata is replicated as part of the image. (This import process is not the same as the import process available in the Catalog utility.)

Table 27-2 Auto Image Replication process overview (*continued*)

Event	Domain in which event occurs	Event description
4	Target primary server (Domain 2)	After the image is imported into the target domain, NetBackup continues to manage the copies in that domain. Depending on the configuration, the media server in Domain 2 can replicate the images to a media server in Domain 3.

One-to-many Auto Image Replication model

In this configuration, all copies are made in parallel. The copies are made within the context of one NetBackup job and simultaneously within the originating storage server context. If one target storage server fails, the entire job fails and is retried later.

All copies have the same **Target Retention**. To achieve different **Target Retention** settings in each target primary server domain, either create multiple source copies or cascade duplication to target primary servers.


Cascading Auto Image Replication model

Replications can be cascaded from the originating domain to multiple domains. Storage lifecycle policies are set up in each domain to anticipate the originating image, import it and then replicate it to the next target primary.

Figure 27-1 represents the following cascading configuration across three domains.

- The image is created in Domain 1, and then replicated to the target Domain 2.
- The image is imported in Domain 2, and then replicated to a target Domain 3.
- The image is then imported into Domain 3.

Figure 27-1 Cascading Auto Image Replication



In the cascading model, the originating primary server for Domain 2 and Domain 3 is the primary server in Domain 1.

Note: When the image is replicated in Domain 3, the replication notification event indicates that the primary server in Domain 2 is the originating primary server. However, after the image is imported successfully into Domain 3, NetBackup correctly indicates that the originating primary server is in Domain 1.

The cascading model presents a special case for the Import SLP that replicates the imported copy to a target primary. (This primary server that is neither the first nor the last in the string of target primary servers.)

The Import SLP must include at least one operation that uses a **Fixed** retention type and at least one operation that uses a **Target Retention** type. So that the Import SLP can satisfy these requirements, the import operation must use a **Target Retention**.

Table 27-3 shows the difference in the import operation setup.

Table 27-3 Import operation difference in an SLP configured to replicate the imported copy


Import operation criteria	Import operation in a cascading model
The first operation must be an import operation.	Same; no difference.
A replication to target primary must use a Fixed retention type	Same; no difference.
At least one operation must use the Target retention .	Here is the difference: To meet the criteria, the import operation must use Target retention .

The target retention is embedded in the source image.

In the cascading model that is represented in [Figure 27-1](#), all copies have the same **Target Retention**—the **Target Retention** indicated in Domain 1.

For the copy in Domain 3 to have a different target retention, add an intermediary replication operation to the Domain 2 storage lifecycle policy. The intermediary replication operation acts as the source for the replication to target primary. Since the target retention is embedded in the source image, the copy in Domain 3 honors the retention level that is set for the intermediary replication operation.

Figure 27-2 Cascading replications to target primary servers, with various target retentions



About the domain relationship for replication

The following table describes important configuration differences depending on the devices that are used for NetBackup Auto Image Replication.

Table 27-4 Replication relationships

Storage	Domain Relationship
NetBackup managed storage	For NetBackup managed storage, establish the relationship between the originating domain and the target domain or domains by setting the properties in the source storage server. NetBackup managed storage that qualifies for replication is Media Server Deduplication Pool storage. See the NetBackup Deduplication Guide. Before you configure the replication relationship, you can add the target primary server as a trusted host. See “About trusted primary servers for Auto Image Replication” on page 974.

Table 27-4 Replication relationships (*continued*)

Storage	Domain Relationship
Third-party vendor disk appliance	For the third-party vendor appliance storage that is exposed through the OpenStorage API, the disk appliance manages the storage. The relationship between the originating domain and the target domain or domains is established by using the storage vendor's tools. The originating NetBackup domain has no knowledge of the storage server in the target domain or domains. When the appliances are configured properly, NetBackup images on the originating disk appliance are replicated automatically to the target disk appliance. That disk appliance uses the OpenStorage API to notify NetBackup that a replication event occurred. NetBackup then imports those images. NetBackup manages the lifecycle of the backup images but does not manage the storage. Configuring the disk appliance replication relationship is beyond the scope of the NetBackup documentation.

Caution: Choose the target storage server carefully. A target storage server must not also be a storage server for the originating domain.

About the replication topology for Auto Image Replication

For Auto Image Replication, the disk volumes have the properties that define the replication relationships between the volumes. The knowledge of the volume properties is considered the replication topology. The following are the replication properties that a volume can have:

- Source** A source volume contains the backups of your clients. The volume is the source for the images that are replicated to a remote NetBackup domain. Each source volume in an originating domain has one or more replication partner target volumes in a target domain.
- Target** A target volume in the remote domain is the replication partner of a source volume in the originating domain.
- None** The volume does not have a replication attribute.

The following describes the replication topology for the supported storage types:

Table 27-5 Replication topology by storage type

Storage type	Replication topology
Media Server Deduplication Pool	NetBackup exposes the storage for a Media Server Deduplication Pool as a single volume. Therefore, there is always a one-to-one volume relationship for MSDP. You configure the replication relationships in the source domain. See the NetBackup Deduplication Guide.
Disk storage devices that support replication and also conform to the Cohesity OpenStorage API	Your storage administrator configures the replication topology of the volumes in the storage devices. Based on the volume properties, you create homogeneous disk pools. That is, all of the volumes in a disk pool must have the same properties, and you create the disk pools to match that topology. The disk pools inherit the replication properties from the volumes that you add to them. You should work with your storage administrator to understand the topology so you can create the proper disk pools. You should also work with your storage administrator to understand any changes that are made to the replication topology. NetBackup discovers the topology of the volumes when you configure a disk pool.

NetBackup discovers the replication topology when you configure the replication relationships. NetBackup discovers topology changes when you use the **Refresh** option in **Storage > Disk storage > Disk pools**.

NetBackup includes a command that can help you understand your replication topology. Use the command in the following situations:

- After you configure the replication targets.
- After you configure the storage server and before you configure disk pools.
- After changes to the volumes that comprise the storage.

See [“Viewing the replication topology for Auto Image Replication”](#) on page 968.

Viewing the replication topology for Auto Image Replication

A volume that is a source of replication must have at least one replication partner that is the target of the replication. NetBackup lets you view the replication topology of the storage.

See [“About the replication topology for Auto Image Replication”](#) on page 967.

To view the replication topology for Auto Image Replication

- ◆ Run the `bpstsinfo` command, specifying the storage server name and the server type. The following is the command syntax:
 - **Windows:** `install_path\NetBackup\bin\admincmd\bpstsinfo -lsuinfo -storage_server host_name -stype server_type`
 - **UNIX:** `/usr/opensv/netbackup/bin/admincmd/bpstsinfo -lsuinfo -storage_server host_name -stype server_type`

The following are the options and arguments for the command:

<code>-storage_server host_name</code>	The name of the target storage server.
<code>-stype server_type</code>	■ For a Media Server Deduplication Pool target, use <code>PureDisk</code>. ■ For an OpenStorage disk appliance, the vendor provides the string for <code>server_type</code>.

Save the output to a file so that you can compare the current topology with the previous topology to determine what has changed.

See [“Sample volume properties output for MSDP replication”](#) on page 969.

See [“Sample volume properties output for OpenStorage backup replication”](#) on page 971.

Sample volume properties output for MSDP replication

The following two examples show output from the `bpstsinfo -lsuinfo` command for two NetBackup deduplication storage servers. The first example is the output from the source disk pool in the originating domain. The second example is from the target disk pool in the remote primary server domain.

The two examples show the following:

- All of the storage in a deduplication disk pool is exposed as one volume: `PureDiskVolume`.
- The `PureDiskVolume` of the deduplication storage server `bit1.datacenter.example.com` is the source for the replication operation.
- The `PureDiskVolume` of the deduplication storage server `target_host.dr-site.example.com` is the target of the replication operation.

```
> bpstsinfo -lsuinfo -storage_server bit1.datacenter.example.com -stype PureDisk
LSU Info:
```

```

Server Name: PureDisk:bit1.datacenter.example.com
LSU Name: PureDiskVolume
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: PureDisk storage unit (/bit1.datacenter.example.com#1/2)
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE | STS_LSUF_STORAGE_NOT_FREED |
      STS_LSUF_REP_ENABLED | STS_LSUF_REP_SOURCE)
Save As : (STS_SA_CLEARF | STS_SA_IMAGE | STS_SA_OPAQUEF)
Replication Sources: 0 ( )
Replication Targets: 1 ( PureDisk:target_host.dr-site.example.com:PureDiskVolume )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 74645270666
Physical Size: 77304328192
Bytes Used: 138
Physical Bytes Used: 2659057664
Resident Images: 0

```

```

> bpstsinfo -lsuinfo -storage_server target_host.dr-site.example.com -stype PureDisk
LSU Info:

```

```

Server Name: PureDisk:target_host.dr-site.example.com
LSU Name: PureDiskVolume
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: PureDisk storage unit (/target_host.dr-site.example.com#1/2)
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE | STS_LSUF_STORAGE_NOT_FREED |
      STS_LSUF_REP_ENABLED | STS_LSUF_REP_TARGET)
Save As : (STS_SA_CLEARF | STS_SA_IMAGE | STS_SA_OPAQUEF)
Replication Sources: 1 ( PureDisk:bit1:PureDiskVolume )
Replication Targets: 0 ( )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 79808086154
Physical Size: 98944983040
Bytes Used: 138
Physical Bytes Used: 19136897024
Resident Images: 0

```

Sample volume properties output for OpenStorage backup replication

The following examples show sample output from the `bpstsinfo` command for two OpenStorage devices. The first example is the output from the source disk pool that contains the client backups. The second example is from the target disk pool in the remote primary server domain.

The two examples show the following:

- Volume `dv01` on storage server `pan1.example.com` is the replication source for volume `dv01` on `pan2.example.com`.
- Volume `dv02` on storage server `pan1.example.com` is the replication source for volume `dv02` on `pan2.example.com`.
- Volume `dv03` on both devices has no replication properties.

```
>bpstsinfo -lsuinfo -storage_server pan1.example.com -stype Pan
```

LSU Info:

```
Server Name: pan1.example.com
LSU Name: dv01
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
      STS_LSUF_REP_SOURCE)
Save As : (STS_SA_IMAGE)
Replication Sources: 0 ( )
Replication Targets: 1 ( Pan:pan2.example.com:dv01 )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 80525455360
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0
```

LSU Info:

```
Server Name: pan1.example.com
LSU Name: dv02
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
```

```
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
STS_LSUF_REP_SOURCE)
Save As : (STS_SA_IMAGE)
Replication Sources: 0 ( )
Replication Targets: 1 ( Pan:pan2.example.com:dv02 )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 80525455360
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0
```

LSU Info:

```
Server Name: pan1.example.com
LSU Name: dv03
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED)
Save As : (STS_SA_IMAGE)
Replication Sources: 0 ( )
Replication Targets: 0 ( )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 80525455360
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0
```

>bpstsinfo -lsuinfo -storage_server pan2.example.com -stype Pan

LSU Info:

```
Server Name: pan2.example.com
LSU Name: dv01
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
STS_LSUF_REP_TARGET)
```


Save As : (STS_SA_IMAGE)
Replication Sources: 1 (Pan:pan1.example.com:dv01)
Replication Targets: 0 ()
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 80525455360
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0

LSU Info:

Server Name: pan2.example.com
LSU Name: dv02
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
STS_LSUF_REP_TARGET)
Save As : (STS_SA_IMAGE)
Replication Sources: 1 (Pan:pan1.example.com:dv02)
Replication Targets: 0 ()
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 80525455360
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0

LSU Info:

Server Name: pan2.example.com
LSU Name: dv03
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED)
Save As : (STS_SA_IMAGE)
Replication Sources: 0 ()
Replication Targets: 0 ()
Maximum Transfer: 2147483647

Block Size: 512
 Allocation Size: 0
 Size: 80525455360
 Physical Size: 0
 Bytes Used: 2285355008
 Physical Bytes Used: 0
 Resident Images: 0

About trusted primary servers for Auto Image Replication

NetBackup provides the ability to establish a trust relationship between replication domains. A trust relationship is optional for the Media Server Deduplication Pool as a target storage. Before you configure a storage server as a target storage, establish a trust relationship between the source A.I.R. and the target A.I.R. operations.

The following items describe how a trust relationship affects Auto Image Replication:

No trust relationship	NetBackup replicates to all defined target storage servers. You cannot select a specific host or hosts as a target.
Trust relationship	You can select a subset of your trusted domains as a target for replication. NetBackup then replicates to the specified domains only rather than to all configured replication targets. This type of Auto Image Replication is known as targeted A.I.R.

About adding a trusted primary server using NetBackup CA-signed certificate

With targeted A.I.R., when trust is established between the source and the remote target server, you need to establish trust in both the domains.

1. In the source primary server, add the target primary server as a trusted server.
2. In the target primary server, add the source primary server as a trusted server.

Note: The **NetBackup web UI** does not support adding a trusted primary server using an external CA-signed certificate.

See [“Add a trusted primary server”](#) on page 164.

See [“About the certificate to use to add a trusted primary server”](#) on page 163.

The following diagram illustrates the different tasks for adding trusted primary servers when NetBackup CA-signed certificate (or host ID-based certificate) is used to establish trust between the source and the target primary servers.

Figure 27-3 Tasks to establish a trust relationship between primary servers for targeted A.I.R. using NetBackup CA-signed certificate




Table 27-6 Tasks to establish a trust relationship between primary servers for targeted A.I.R.

Step	Task	Procedure
Step 1	Administrators of both the source and the target primary servers must obtain each other's CA certificate fingerprint and authorization tokens or the user credentials. This activity must be performed offline. Note: It is recommended to use an authentication token to connect to the remote primary server. An authentication token provides restricted access and allows secure communication between both the hosts. The use of user credentials (user name and password) may present a possible security breach.	To obtain the authorization tokens, use the <code>bpnbat</code> command to log on and <code>nbcertcmd</code> to get the authorization tokens. To obtain the SHA1 fingerprint of root certificate, use the <code>nbcertcmd -displayCACertDetail</code> command. To perform this task, see the NetBackup Commands Reference Guide. Note: When you run the commands, keep the target as the remote server.

Table 27-6 Tasks to establish a trust relationship between primary servers for targeted A.I.R. (*continued*)

Step	Task	Procedure
Step 2	Establish trust between the source and the target domains. - On the source primary server, add the target primary server as trusted server. - On the target primary server, add the source primary server as trusted server.	To perform this task in the NetBackup web UI, see the following topic: See “Add a trusted primary server” on page 164. To perform this task using the <code>nbseccmd</code>, see the NetBackup Commands Reference Guide.
Step 3	After you have added the source and target trusted servers, they have each other’s host ID-based certificates. The certificates are used during each communication. Primary Server A has a certificate that Primary Server B issued and vice versa. Before communication can occur, Primary Server A presents the certificate that Primary Server B issued and vice versa. The communication between the source and the target primary servers is now secured.	To understand the use of host ID-based certificates, see the NetBackup Security and Encryption Guide.
Step 3.1	Configure the source media server to get the security certificates and the host ID certificates from the target primary server.	See the <i>NetBackup Deduplication Guide</i>. http://www.veritas.com/docs/DOC5332
Step 4	Create an import storage lifecycle policy in the target domain. Note: The import storage lifecycle policy name should contain less than or equal to 112 characters.	See “About storage lifecycle policies” on page 601.
Step 5	On the source MSDP server, add the credentials of the target storage server.	See the <i>NetBackup Deduplication Guide</i>. http://www.veritas.com/docs/DOC5332
Step 5.1	Create a replication storage lifecycle policy in the source domain using the specific target primary server and storage lifecycle policy. The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains.	See “About storage lifecycle policies” on page 601.
Step 6	The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as Auto Image Replication.	See “About NetBackup Auto Image Replication” on page 961.

About the storage lifecycle policies required for Auto Image Replication

To replicate images from one NetBackup domain to another NetBackup domain requires two storage lifecycle policies. The following table describes the policies and their requirements:

Table 27-7 SLP requirements for Auto Image Replication

Domain	Storage lifecycle policy requirements
Domain 1 (Source domain)	The Auto Image Replication SLP in the source domain must meet the following criteria: ■ The first operation must be a Backup operation to storage that NetBackup supports for replication. Indicate the exact storage unit from the drop-down list. Do not select Any Available. Note: The target domain must contain the same type of storage to import the image. ■ At least one operation must be a Replication operation to storage in another domain that NetBackup supports for replication from the source storage. You can configure multiple Replication operations in an Auto Image Replication SLP. The Replication operation settings determine whether the backup is replicated to all replication targets in all primary server domains or only to specific replication targets. ■ The SLP must be of the same data classification as the Import SLP in Domain 2.
Domain 2 (Target domain)	If replicating to all targets in all domains, in each domain NetBackup automatically creates an Import SLP that meets all the necessary criteria. Note: If replicating to specific targets, you must create the Import SLP before creating the Auto Image Replication SLP in the originating domain. The Import SLP must meet the following criteria: ■ The first operation in the SLP must be an Import operation. NetBackup must support the Destination storage as a target for replication from the source storage. Indicate the exact storage unit from the drop-down list. Do not select Any Available. ■ The SLP must contain at least one operation that has the Target retention specified. ■ The SLP must be of the same data classification as the SLP in Domain 1. Matching the data classification keeps a consistent meaning to the classification and facilitates global reporting by data classification.

Figure 27-4 shows how the SLP in the target domain is set up to replicate the images from the originating primary server domain.

Figure 27-4 Storage lifecycle policy pair required for Auto Image Replication

Edit Storage Lifecycle Policy

Storage lifecycle policy

Validation report

Storage lifecycle policy name

SLP-MSDP-Rep

Data classification

No data classification

Priority for secondary operations

0

A higher number is greater priority.

+ Add

	Window	Storage	Volume pool	Media owner	Retention type	Retention period
☐	Operation					
☐	Backup		stu_local		Fixed	2 weeks
☐	Replication	Default_24x7_Window	SLP-MSDP-Rep		Fixed	2 weeks

2 Records

State of secondary operation processing

To find impact on policies associated with this SLP due to change in configuration click here.

CancelSave

Edit Storage Lifecycle Policy

Storage lifecycle policy

Validation report

Storage lifecycle policy name

SLP-MSDP-Rep

Data classification

No data classification

Priority for secondary operations

0

A higher number is greater priority.

+ Add

	Window	Target primary	Storage	Storage type	Volume pool	Media owner
☐	Import	Default_24x7_Window	stu_local_sadto6vm00	PureDisk		

1 Records

State of secondary operation processing

To find impact on policies associated with this SLP due to change in configuration click here.

CancelSave

Note: Restart `nbtstserv` after you make changes to the underlying storage for any operation in an SLP.

Replicating to all inter-domain targets or to a specific target

With Auto Image Replication, you can replicate backups to all configured replication targets or to a subset of all configured replication targets. To replicate to specific primary server domains, you must first configure trusted primary servers.

See [“About trusted primary servers for Auto Image Replication”](#) on page 974.

Table 27-8 Configuration differences between replicating to all inter-domain targets or to a specific target

Replication objective	Auto Image Replication SLP configuration	Import SLP configuration
Replicate the backup to all configured primary server domains.	Create an SLP in the originating domain. ■ The first operation must be a Backup operation. ■ The SLP must include a Replication operation. To copy to all domains, select All inter-domain replication target(s).	The Import SLPs are created automatically in all domains. See Figure 27-5 for a representation of this scenario.
Replicate the backup to targets in specific NetBackup primary server domains.	In this case, first create the Import SLPs before the SLP in the originating domain. ■ The first operation must be a Backup operation. ■ The SLP must include a Replication operation. Select A specific Master Server and indicate the domain of the target primary server.	The Import SLPs are not created automatically. Note: Create the Import SLP before creating the Auto Image Replication SLP in the originating domain. Create an Import SLP in each target domain. The Import SLP must have an Import operation as the first operation, but can contain other operations as needed. Figure 27-6 represents this scenario.

Additional requirements for Auto Image Replication SLPs are described in the following topic:

See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 977.

Figure 27-5 Replicating from one domain to all inter-domain primary servers





Figure 27-6 represents replication to a specific primary domain target.

Figure 27-6 Replicating from one domain to a specific inter-domain primary server



To replicate copies to a specific domain, make sure that the target domain primary server is a trusted primary server of the originating domain.

See [“Add a trusted primary server”](#) on page 164.

Customizing how nbstserv runs duplication and import jobs

The NetBackup Storage Lifecycle Manager (`nbstserv`) runs replication, duplication, and import jobs. Both the Storage Lifecycle Manager service and the Import Manager service run within `nbstserv`.

The NetBackup administrator can customize how `nbstserv` runs jobs by changing the default of certain SLP-related configuration options.

See [“SLP settings properties”](#) on page 168.

About Auto Image Replication import confirmation

When using targeted Auto Image Replication (A.I.R.), storage lifecycle (SLP) processing of each replicated image is paused in the source domain until a message has been received from the target domain that confirms that the image has been imported successfully. SLP processing does not continue for the image until the confirmation occurs. Import confirmation ensures that source domain images remain in place at least until those images have been safely imported in the target domain.

In the source domain, NetBackup keeps track of image copies that remain in import pending state for more than 24 hours. Whenever such images exist, NetBackup generates a message in the **Problems** report. This message indicates that import-pending images are present. You can configure a different time threshold for generating a confirmation message. You can also configure an email address to receive the same information send to the **Problems** report.

See [“SLP settings properties”](#) on page 168.

If you receive notification that images are in an import pending state, you can run the `nbstlutil` command. This command generates a report that provides details of those images. You must then investigate the cause of the import problem and address any issue that you find. The import-pending state is automatically updated when the imports are completely successfully. Alternatively, you may decide that import operations in the target domain are no longer needed, and you can then cancel SLP processing for those images in the source domain. This action also clears the import-pending state, but no further SLP processing happens for those images.

See the [NetBackup Commands Reference Guide](#) for details about `nbstlutil`.

Auto Image Replication setup overview

The following table is an overview of the setup process for Auto Image Replication, describing the actions that are required.

Table 27-9 Auto Image Replication setup overview

Step	Action	Description
1	Configure the storage servers	Configure the storage servers for your storage type. See the NetBackup Deduplication Guide or the NetBackup OpenStorage Solutions Guide for Disk .
2	Configure the disk pools.	Configure the disk pools for your storage type. To replicate images from one domain to another requires that suitable storage be configured in each domain. NetBackup must support the storage for replication. See the NetBackup Deduplication Guide or the NetBackup OpenStorage Solutions Guide for Disk .
3	Configure the storage units.	Configure the storage units in both the originating domain and the target domain. See “Creating a storage unit” on page 548.
4	Define the relationship between the domains.	Define the relationship between the domains so that the originating domain knows where to send the data. See “About the domain relationship for replication” on page 966. See “About trusted primary servers for Auto Image Replication” on page 974.
5	Configure the storage lifecycle policies.	Configure the necessary storage lifecycle policies. See “About the storage lifecycle policies required for Auto Image Replication” on page 977. See “Creating a storage lifecycle policy” on page 602.
6	Configure and run the backup policy in the originating domain.	The backup policy must indicate the configured SLP as the Policy storage selection. See “Creating a backup policy” on page 657.

How to resolve volume changes for Auto Image Replication

When you edit a disk pool, NetBackup loads the disk pool properties from the catalog. NetBackup queries the storage server for changes when you either click the **Refresh** button on the **Disk pools** tab or when you select the **Update disk volume** option.

It is recommended that you take the following actions when the volume topology changes:

- Discuss the changes with the storage administrator. You need to understand the changes so you can change your disk pools (if required) so that NetBackup can continue to use them.
- If the changes were not planned for NetBackup, ask your storage administrator to revert the changes so that NetBackup functions correctly again.

NetBackup can process changes to the following volume properties:

- Replication source
- Replication target
- None

If these volume properties change, NetBackup can update the disk pool to match the changes. NetBackup can continue to use the disk pool, although the disk pool may no longer match the storage unit or storage lifecycle purpose.

The following table describes the possible outcomes and how to resolve them.

Table 27-10 Refresh outcomes

Outcome	Description
No changes are discovered.	No changes are required.
NetBackup discovers the new volumes that you can add to the disk pool.	The new volumes display in the properties for that Disk pool.

Table 27-10 Refresh outcomes (*continued*)


Outcome	Description
The replication properties of all of the volumes changed, but they are still consistent.	A Disk Pool Configuration Alert pop-up box notifies you that the properties of all of the volumes in the disk pool changed, but they are all the same (homogeneous).  You must click OK in the alert box, after which the disk pool properties in the Change Disk Pool dialog box are updated to match the new volume properties If new volumes are available that match the new properties, NetBackup displays those volumes in the Change Disk Pool dialog box. You can add those new volumes to the disk pool. In the Change Disk Pool dialog box, select one of the following two choices: ■ OK. To accept the disk pool changes, click OK in the Change Disk Pool dialog box. NetBackup saves the new properties of the disk pool. NetBackup can use the disk pool, but it may no longer match the intended purpose of the storage unit or storage lifecycle policy. Change the storage lifecycle policy definitions to ensure that the replication operations use the correct source and target disk pools, storage units, and storage unit groups. Alternatively, work with your storage administrator to change the volume properties back to their original values. ■ Cancel. To discard the changes, click Cancel in the Change Disk Pool dialog box. NetBackup does not save the new disk pool properties. NetBackup can use the disk pool, but it may no longer match the intended use of the storage unit or storage lifecycle policy.

Table 27-10 Refresh outcomes (*continued*)



Outcome	Description
The replication properties of the volumes changed, and they are now inconsistent.	A Disk Pool Configuration Error pop-up box notifies you that the replication properties of some of the volumes in the disk pool changed. The properties of the volumes in the disk pool are not homogeneous.  You must click OK in the alert box. In the Change Disk Pool dialog box, the properties of the disk pool are unchanged, and you cannot select them (that is, they are dimmed). However, the properties of the individual volumes are updated. Because the volume properties are not homogeneous, NetBackup cannot use the disk pool until the storage configuration is fixed. NetBackup does not display new volumes (if available) because the volumes already in the disk pool are not homogeneous. To determine what has changed, compare the disk pool properties to the volume properties. See “Viewing the replication topology for Auto Image Replication” on page 968. Work with your storage administrator to understand the changes and why they were made. The replication relationships may or may not have to be re-established. If the relationship was removed in error, re-establishing the relationships seem justified. If you are retiring or replacing the target replication device, you probably do not want to re-establish the relationships. The disk pool remains unusable until the properties of the volumes in the disk pool are homogenous. In the Change Disk Pool dialog box, click OK or Cancel to exit the Change Disk Pool dialog box.

Table 27-10 Refresh outcomes (continued)

Outcome	Description
NetBackup cannot find a volume or volumes that were in the disk pool.	A Disk Pool Configuration Alert pop-up box notifies you that an existing volume or volumes was deleted from the storage device:  NetBackup can use the disk pool, but data may be lost. To protect against accidental data loss, NetBackup does not allow volumes to be deleted from a disk pool. To continue to use the disk pool, do the following: ■ Use the <code>bpimmedia</code> command or the Images on Disk report to display the images on the specific volume. ■ Expire the images on the volume. ■ Use the <code>nbdevconfig</code> command to set the volume state to DOWN so NetBackup does not try to use it.

Removing or replacing replication relationships in an Auto Image Replication configuration

Auto Image Replication replicates backups from a storage server in the source domain to storage servers in one or more target NetBackup domains. If a storage server needs to be removed or replaced from such a relationship, the involved domains need to make proper preparations to stop or to redirect replication. That is, to remove the replication relationship.

For example, consider a scenario where cascading Auto Image Replication is used across 3 domains. To remove the storage server in domain 2 (which serves as the destination of the import from Domain 1) preparations need to be made in all three domains. The preparations include modifying storage lifecycle policies and removing the storage server from the replication topology.

Figure 27-7 Removing target storage server example




Table 27-11 contains topics that address processes that are involved in removing or replacing relationships in an Auto Image Replication configuration:

Table 27-11

Topic	Reference
Adding a replication relationship between two storage servers.	See “Adding or removing a replication relationship between two storage servers” on page 988.
Removing a replication relationship between two storage servers.	
Removing a replication relationship between a domain and a storage server.	See “Removing all replication relationships between a domain and a storage server” on page 988.
Replacing a replication relationship between a domain and a storage server.	See “Replacing all replication relationships between a domain and a storage server” on page 989.
Removing all replication relationships involving a storage server.	See “Removing or replacing all replication relationships involving a storage server” on page 993.
Replacing all replication relationships involving a storage server.	

Adding or removing a replication relationship between two storage servers

Before adding a replication relationship, you must understand the consequences of this action:

- Any classic (or non-targeted) Auto Image Replication storage lifecycle policies that replicate from the source domain storage server will replicate to the target domain storage server.

Before removing a replication relationship, you must understand the consequences of this action:

- Any classic (or non-targeted) A.I.R. SLPs that replicate from the source domain storage server will no longer replicate to the target domain storage server.
- Any targeted A.I.R. SLPs that replicate from the source domain storage server to the target domain storage server will fail. These SLPs should be removed or modified so that they no longer replicate to the target domain.

Complete the following steps to add or remove a replication relationship between two storage servers:

To change (add or remove) a replication relationship

- 1 This step depends on whether OpenStorage or MSDP storage is used:
 - For OpenStorage storage, contact your storage administrator to change the replication relationship, as this procedure differs between vendors.
 - For MSDP storage, the relationships can be changed in the source domain by the NetBackup administrator.
- 2 After the relationship is changed, update the disk pools in both domains to reflect the topology changes:

Select **Storage > Disk storage**. Then select the **Disk pools** tab. Select and refresh the disk pools. The `nbdevconfig -updatedp` command can also be used. (See the *OpenStorage Solutions Guide* for information about the replication topology for Auto Image Replication.)

Removing all replication relationships between a domain and a storage server

Complete the following steps to remove all replication relationships between two domains. The procedure refers to domain D1 and a storage server in domain D2:

To remove all replication relationships between a domain and a storage server

- 1 In domain D1:

Run the following command on the primary server:


```
nbdecommission -list_ref -oldserver  
storage_server_name-machinetype replication_host
```

The output lists:

- All of the storage servers in the domain that reference the deprecated storage server as a replication target or a replication source.
- All of the targeted A.I.R. replication SLPs that reference an SLP that imports to the deprecated storage server.

Note: If an SLP has in-process images, either wait until those images are complete or cancel them before decommissioning the storage server. Note that this includes SLPs of all versions. Use the SLP utility command (`nbstlutil`) to cancel the processing of existing SLP-managed images.

See [“Lifecycle operation administration using the nbstlutil command”](#) on page 605.

- 2 Remove the replication operations from the targeted A.I.R. storage lifecycle polices that were found in step 1.

If these SLPs are no longer necessary, the SLPs can be deleted now.

- 3 Remove the replication relationships that were found in step 1.

See [“Adding or removing a replication relationship between two storage servers”](#) on page 988.

Any remaining A.I.R. storage lifecycle polices in both domains that are no longer necessary can be deleted now.

- 4 Run the following command to decommission the storage server:

```
nbdecommission -oldserver storage_server_name-machinetype  
replication_host
```

Replacing all replication relationships between a domain and a storage server

Note: If the storage server to be replaced is involved in a classic (non-targeted) Auto Image Replication configuration, the replacement storage server *must not* have any additional replication target relationships.

The procedure refers to two domains: D1 and D2. Auto Image Replication is configured to occur from storage servers in D1 to D2. Also, from D2 to D1.

A new storage server (S2) is added to D2 so that D2 now contains two storage servers (S1 and S2). Relationships to S1 must be replaced with relationships to S2.

[Table 27-12](#) lists the configuration before the switch to D2. [Table 27-13](#) lists the configuration after the changes have been made.

Table 27-12 Example configuration before changes

Domains	Storage servers	Storage lifecycle policies
D1	Contains several storage servers but they are not necessary in the example.	■ BACKUP_D1_REPLICATE_D2 This SLP replicates to S1 in D2 using the target import SLP IMPORT_S1. ■ IMPORT_D1 This SLP imports to a storage server in D1.
D2	S1 (To be deprecated) S2 (New in D2)	■ BACKUP_D2_REPLICATE_D1 This SLP replicates to D1 using the target import SLP IMPORT_D1. ■ IMPORT_S1 This SLP imports to D2 and stores on S1.

Figure 27-8 Topology before changes





Table 27-13 Example configuration after changes

Domains	Storage servers	Storage lifecycle policies
D1	Contains several storage servers but they are not necessary in the example.	■ BACKUP_D1_REPLICATE_D2 This SLP replicates to S2 in D2 using the target import SLP IMPORT_S2 ■ IMPORT_S1 This SLP imports to a storage server in D1.
D2	S1 (To be deprecated) S2 (New in D2)	■ BACKUP_D2_REPLICATE_D1 This SLP replicates to D1 using the target import SLP IMPORT_D1. ■ IMPORT_S2 This SLP imports to D2 and stores on S2.

Figure 27-9 Topology after changes



Circled items have been changed.

In the following procedure, note that steps must be performed in specific domains.

Table 27-14 Replacing all replication relationships between a domain and a storage server

Domain	Step	Description
In domain D2:	1	Run the following command to determine the replication relationships of storage server S1: <pre>bpstsinfo -lsuinfo -storage_server storage_server_name -stype storage_server_type</pre> For example: <pre>bpstsinfo -lsuinfo -storage_server S1 -stype replication_host</pre>
	2	For each replication target relationship that is found in step 1, add a corresponding relationship between storage server S2 and the target in D1. See “Adding or removing a replication relationship between two storage servers” on page 988.
	3	Copy all import SLPs that import to storage server S1. Modify the new SLPs to import to storage server S2. For example, copy IMPORT_S1. Modify the SLP to import to storage server S2. Change the SLP name to reflect that it will import to S2: IMPORT_S2. Note: Do not delete the existing import SLPs yet. All of the SLPs that are no longer performing any function can be deleted later in this procedure.
In domain D1:	4	Run the following command on the primary server: <pre>nbdecommission -list_ref -oldserver storage_server_name-machinetype replication_host</pre> For example: <pre>nbdecommission -list_ref -oldserver S1 -machinetype replication_host</pre> The output lists: ■ All of the storage servers in the source domain (S1) that reference the deprecated storage server as a replication target or replication source. ■ All of the targeted A.I.R. replication SLPs that reference an SLP that imports to the deprecated storage server. Note: If an SLP has in-process images, either wait until those images are complete or cancel them before decommissioning the storage server. Note that this includes SLPs of all versions. Use the SLP utility command (<code>nbstlutil</code>) to cancel the processing of existing SLP-managed images. See “Lifecycle operation administration using the nbstlutil command” on page 605.

Table 27-14 Replacing all replication relationships between a domain and a storage server (*continued*)

Domain	Step	Description
In domain D2:		Complete steps 5-6 only if step 4 listed any replication source relationships.
	5	Modify replication SLPs that replicate from S1 to domain D1 to replicate from S2. This includes both non-targeted and targeted A.I.R. SLPs. For example, change the Replication operation in BACKUP_D2_REPLICATE_D1. Change the backup storage destination from S1 to S2.
	6	Remove the replication source relationships that were found in step 4. See “Adding or removing a replication relationship between two storage servers” on page 988.
In domain D1:		Complete steps 7-10 only if step 4 listed any replication target relationships.
	7	Add replication relationships from each storage server that was listed in step 2 that has a replication target relationship to storage server S2. See “Adding or removing a replication relationship between two storage servers” on page 988.
	8	Modify the replication operations that were found in step 4. Change the target import SLP to the corresponding import SLP created in step 3. For example, change the Replication operation in BACKUP_D1_REPLICATE_D2. Change the Target import SLP setting from IMPORT_S1 to IMPORT_S2.
	9	Remove the replication target relationships that were found in step 4. See “Adding or removing a replication relationship between two storage servers” on page 988.
	10	Run the following command to decommission the storage server: <pre>nbdecommission -oldserver storage_server_name -machinetype replication_host</pre> For example: <pre>nbdecommission -oldserver D1 -machinetype replication_host</pre>
In domain D2:	11	Any import SLPs to S1 that are no longer necessary may now be deleted. For example, IMPORT_S1 can be deleted now.

Removing or replacing all replication relationships involving a storage server

To remove or replace all of the replication relationships involving a storage server, completely remove or replace all replication relationships between a domain and

the storage server. This must be done for each domain that is involved in a replication relationship with the storage server.

The following command lists replication target and replication source relationships. The command is useful for determining which domains have replication relationships with the storage server:

```
bpstsinfo -lsuinfo -storage_server storage_server_name -stype
storage_server_type
```

Example: Replacing a storage server in a non-targeted Auto Image Replication configuration

This example walks through the steps necessary to replace a storage server in a simple, classic (non-targeted) Auto Image Replication configuration.

Specifically, to replace an MSDP storage server (D2_MSDP_1) in domain D2 with another MSDP storage server (D2_MSDP_2).

Table 27-15 Example configuration

Domains	Storage servers	Storage lifecycle policies
D1	D1_MSDP	BACKUP_D1
D2	D2_MSDP_1 D2_MSDP_2	No replication-related SLPs

The changes to the replication and the storage lifecycle topologies are tracked throughout the example.

Figure 27-10 Example replication topology before the process






Figure 27-11 Example storage lifecycle policy topology




To replace the storage server D2_MSDP_1 with D2_MSDP_2

- 1** Add a replication target relationship from D1_MSDP to D2_MSDP_2.
- 2** Refresh disk pools in both domains.



- 3** Remove the replication target relationship between D1_MSDP and D2_MSDP_1.

4 Refresh disk pools in both domains.



5 In D1, run the following command to decommission MSDP_D2_1:

```
nbdecommission -oldserver MSDP_D2_1 -machinetype replication_host
```

Example: Replacing a storage server in a cascading, targeted Auto Image Replication configuration

This example walks through the steps necessary to replace an MSDP storage server in a cascading, targeted Auto Image Replication configuration. The changes to the replication and the storage lifecycle topologies are tracked throughout the example.

The environment contains three domains. Each domain contains one or more MDSP storage servers.

Table 27-16 Example configuration

Domains	Storage servers	Storage lifecycle policies
D1	D1_MSDP	BACKUP_D1
D2	D2_MSDP D2_MSDP_2	IMPORT_D2
D3	D3_MSDP	IMPORT_D3

Figure 27-12 Example replication topology






Figure 27-13 Example storage lifecycle policy topology




To replace the storage server D2_MSDP with D2_MSDP_2


- 1** Add replication target relationship from D2_MSDP_2 to D3_MSDP.
- 2** Refresh disk pools in both domain D2 and domain D3.




- 3 Copy IMPORT_D2 and modify name to IMPORT_2_D2 and destination storage to a storage unit on D2_MSDP_2.




- 4 Add a replication target relationship from D1_MSDP to D2_MSDP_2.
- 5 Refresh disk pools in both domains.



- 6 Modify BACKUP_D1 Replication Operation Target import SLP to IMPORT_2_D2.




- 7 Remove the replication target relationship from D1_MSDP to D2_MSDP_1.
- 8 Refresh the disk pools in both domains.



- 9 Run the following command in D1:


```
nbdecommission -oldserver MSDP_D2 -machinetype replication_host
```

10 Delete IMPORT_D2.



11 Remove the replication target relationship from D2_MSDP_1 to D3_MSDP.

12 Refresh disk pools in both domains.



About restoring from a backup at a target primary domain

While it is possible to restore a client directly by using the images in the target primary domain, do so only in a disaster recovery situation. In this discussion, a disaster recovery situation is one in which the originating domain no longer exists and clients must be recovered from the target domain

Table 27-17 Client restores in disaster recovery scenarios

Disaster recovery scenario	Does client exist?	Description
Scenario 1	Yes	Configure the client in another domain and restore directly to the client.
Scenario 2	No	Create the client in the recovery domain and restore directly to the client. This is the most likely scenario.
Scenario 3	No	Perform an alternate client restore in the recovery domain.

The steps to recover the client are the same as any other client recovery. The actual steps depend on the client type, the storage type, and whether the recovery is an alternate client restore.

For restores that use Granular Recovery Technology (GRT), an application instance must exist in the recovery domain. The application instance is required so that NetBackup has something to recover to.

For information on granular recovery, see the following topics and guides:

- See [“Active Directory granular backups and recovery”](#) on page 835.
- See [“Enable granular recovery \(policy attribute\)”](#) on page 702.
- See [“Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology \(GRT\)”](#) on page 1144.
- The [NetBackup for Microsoft SharePoint Server Administrator's Guide](#)
- The [NetBackup for Microsoft Exchange Server Administrator's Guide](#)

Reporting on Auto Image Replication jobs

The Activity Monitor displays both the **Replication** job and the **Import** job in a configuration that replicates to a target primary server domain.

Table 27-18 Auto Image Replication jobs in the Activity Monitor

Job type	Description
Replication	The job that replicates a backup image to a target primary displays in the Activity Monitor as a Replication job. The Target Master label displays in the Storage Unit column for this type of job. Similar to other Replication jobs, the job that replicates images to a target primary can work on multiple backup images in one instance. The detailed status for this job contains a list of the backup IDs that were replicated.

Table 27-18 Auto Image Replication jobs in the Activity Monitor *(continued)*

Job type	Description
Import	The job that imports a backup copy into the target primary domain displays in the Activity Monitor as an Import job. An Import job can import multiple copies in one instance. The detailed status for an Import job contains a list of processed backup IDs and a list of failed backup IDs. Note that a successful replication does not confirm that the image was imported at the target primary. If the data classifications are not the same in both domains, the Import job fails and NetBackup does not attempt to import the image again. Failed Import jobs fail with a status 191 and appear in the Problems report when run on the target primary server. The image is expired and deleted during an Image Cleanup job. Note that the originating domain (Domain 1) does not track failed imports.

About NetBackup Replication Director

Replication Director is the implementation of NetBackup OpenStorage-managed snapshots and snapshot replication, where the snapshots are stored on the storage systems of partnering companies. OpenStorage is a Cohesity API that lets NetBackup communicate with the storage implementations that conform to the API. Replication Director uses the functions of the OpenStorage partners to perform the following tasks:

- To share disks so that multiple heterogeneous media servers can access the same disk volume concurrently.
- To balance loads and tune performance. NetBackup balances backup jobs and storage usage among the media servers and disk pools.
- To make full use of disk array capabilities, including fast storage provisioning and almost unlimited storage.
- To use as an alternative to off-site vaulting. Storage replication technology provides an efficient means to send copies of user data (files, applications, databases) to off-site storage as part of a disaster recovery plan.

NetBackup stores snapshots of client data on the volumes that are available to the storage server.

Snapshots represent a point-in-time of primary storage data as captured by the storage hardware. NetBackup can then instruct the storage server to replicate the snapshot from primary volumes to other volumes available to the storage server. The snapshot can be replicated to multiple volumes within the storage server, or to storage outside of the storage server, such as a tape device or other disk storage.

Replication Director can accommodate an assortment of scenarios to meet the specific data protection needs of an organization.

Replication Director offers a single NetBackup interface for end-to-end data protection management for the following tasks:

- Unified policy management.
Use the **NetBackup web UI** as the one, centralized backup infrastructure to manage the lifecycle of all data.
- Snapshot copy management.
Use NetBackup to manage the entire lifecycle of the snapshot. Replication Director uses OpenStorage with a media server to access the storage server volumes. No image can be moved, expired, or deleted from the disk array unless NetBackup instructs the storage server to do so.
The instruction to perform the initial snapshot comes from an operation in a NetBackup storage lifecycle policy (SLP). You can create one SLP that instructs NetBackup to create the initial snapshot, to replicate the snapshot to several locations, and to indicate a different retention period for each of the replications. Additional instructions (or operations) can be included in the SLP that create a backup from the snapshot, index the snapshot, and more.
- Global search and restore.
Recovery is available from any storage device in the environment that is defined to NetBackup. This includes recovery from the primary copy or any replicated copy on disk, or from any duplicated copy on disk or tape.

For more information, see the [NetBackup Replication Director Solutions Guide](#).

Monitoring and reporting

- [Chapter 28. Monitoring NetBackup activity](#)
- [Chapter 29. Reporting in NetBackup](#)
- [Chapter 30. Email notifications](#)

Monitoring NetBackup activity

This chapter includes the following topics:

- [About the Activity Monitor](#)
- [Setting Activity Monitor options](#)
- [About the Jobs tab](#)
- [About the Daemons tab](#)
- [About the Processes tab](#)
- [About the Drives tab](#)
- [About the Error Logs tab](#)
- [About the jobs database](#)
- [About the Device Monitor](#)
- [About media mount errors](#)
- [About pending requests and actions](#)

About the Activity Monitor

Use the Activity Monitor in the **NetBackup Administration Console** to monitor and control the following aspects of NetBackup:

Jobs

See [“About the Jobs tab”](#) on page 1009.

The job details are described in the online Help.

Services or Daemons	See “About the Daemons tab” on page 1013.
Processes	See “About the Processes tab” on page 1024.
Drives	See “About the Drives tab” on page 1030.
Error Logs	See “About the Error Logs tab” on page 1031.

- As long as the Activity Monitor is active in the **NetBackup Administration Console**, the `bpjobjd` daemon supplies the job activity status to the Activity Monitor.
Updates to the Activity Monitor occur as jobs are initiated, updated, and completed. Without a refresh cycle, updates occur instantaneously.
The status bar appears in the **Jobs** tab, at the top of the Activity Monitor **Details** pane.


Note: The **Filter** option on the **View** menu is useful for displaying in Activity Monitor only those jobs with specified characteristics. For example, the jobs that were started before a specific date; jobs in the queued state; jobs with status completion codes within a specified range.

The status bar displays the following information:

- The primary server on which the jobs reside.
- The total number of jobs.
- The number of jobs in each of the job states: Active, Queued, Waiting for Retry, Suspended, Incomplete, and Done.
- The number of jobs currently selected.
- The number of NetBackup services or daemons that run.
- The number of drives and the state of each (Active, Down).

The numbers always reflect the actual number of jobs, even when the filter is used.

Figure 28-1 Activity Monitor



Right-click in pane to view the shortcut menu

Setting Activity Monitor options

The following procedure describes how to set the options for the Activity Monitor in the **NetBackup Administration Console**.

To configure options for the Activity Monitor

- 1 In the **NetBackup Administration Console**, in the menu bar, click **View > Options** and select the **Activity Monitor** tab to access configurable options for the Activity Monitor.
- 2 The following options are available to receive a confirmation warning.

Confirm job deletions Prompts the user with a confirmation dialog box when a job is deleted.

Confirm job cancellations Prompts the user with a confirmation dialog box when a job is canceled.

Confirm stop daemons	Enable to display a confirmation dialog box when a daemon is stopped. To discontinue further delete, cancel, or stop daemon confirmations, enable the In the future, do not show this warning option.
Maximum details windows	Specifies the maximum number of Activity Monitor job details, daemon details, and the process details windows that can be displayed at one time.
Automatically refresh display	Enable to refresh data on the Daemons tab and the Processes tab and the job details elapsed time. Other Jobs tab data refreshes independently of the Auto Refresh setting. Enter the rate (in seconds) at which data refreshes in the Daemons tab and the Processes tab.
Show error logs in Activity Monitor	Enable to view NetBackup error logs in the Error Logs tab in the Activity Monitor. You can customize the view of error logs in Activity Monitor by configuring the following options: ■ Select the option from the drop-down menu to configure Show error logs for. The default selection is Error and above which lists all the errors and critical logs. You can select to view specific logs like critical logs, information logs, errors, and warnings. ■ Select the time window next to Show error logs for last to configure viewing logs in the selected time frame. Default time window is set for last 24 hours. ■ Set the auto-refresh time by specifying the Refresh display every option, in minutes. The default refresh time is set to 15 minutes.
Show error logs in Job details window	Enable viewing error logs in the Job Details window. When this option is checked, after clicking a specific Job ID, you can view the log entries in the Job Details window in a separately generated Error Logs tab. In the Job Details window, the errors and critical logs that are specific to the selected job entry, are shown.

- 3 Click **OK** to close the dialog box and apply any changes.

About the Jobs tab

The **Jobs** tab in the Activity monitor displays all of the jobs that are in process or that have completed for the primary server currently selected. The **Jobs** tab also displays details about the jobs. The job details are described in the online Help.

For some backup jobs, a parent job is used to perform pre- and post-processing. Parent jobs display a dash (-) in the Schedule column.

A parent job runs the start and end notify scripts (`PARENT_START_NOTIFY`, `PARENT_END_NOTIFY`) from the primary server. The scripts are located in the following directory:

On Windows: `install_path\NetBackup\bin`

On UNIX: `/usr/openv/netbackup/bin/`

The role of the parent job is to initiate requested tasks in the form of children jobs.

The tasks vary, depending on the backup environment, as follows.

Table 28-1 Tasks initiated by parent jobs

Task	Description
Snapshot Client	The parent job creates the snapshot, initiates children jobs, and deletes the snapshot when complete. Children jobs are created if the Snapshot Client settings are configured to retain snapshots for Instant Recovery, then copy snapshots to a storage unit. (Snapshots and copy snapshots to a storage unit is selected in the policy Schedule Attributes tab.) Children jobs are not created if the Snapshot Client settings are configured to retain snapshots for Instant Recovery, but to create snapshots only. That is, the snapshot is not backed up to a storage unit, so no children jobs are generated. (Snapshots only is selected in the policy Schedule Attributes tab.)
Bare Metal Restore	The parent job runs <code>brmsavecfg</code>, then initiates the backup as a child job. If multistreaming and BMR are used together, the parent job can start multiple children jobs.
Catalog backups	The parent job for catalog backups works with <code>bpdbm</code> to initiate multiple children backup jobs: ■ A NetBackup database backup ■ A file system backup of the primary server ■ A backup of the BMR database, if necessary

Table 28-1 Tasks initiated by parent jobs (*continued*)

Task	Description
Multiple copies	A multiple copies job produces one parent job and multiple child jobs. Child jobs that are part of a multiple copies parent job cannot be restarted individually. Only the parent job (and subsequently all the children jobs) can be restarted. See “Multiple copies (schedule attribute)” on page 747.
Multiple data streams	The parent job performs stream discovery and initiates children jobs. A parent job does not display a schedule in the Activity Monitor. Instead, a dash (-) appears for the schedule because the parent schedule is not used and the children schedules may be different. The children jobs display the ID of the parent job in the Activity Monitor.
SharePoint	The parent job runs a resolver process during which children jobs are started. This process is similar to the stream discovery for multiple data streams. If multiple data streams are enabled, some children jobs can be split into multiple streams.
Vault	The parent job starts the Vault profile. Then, the Vault profile starts the duplicates as jobs. The duplicates do not appear as children jobs in the Activity Monitor.

Viewing job details in the Activity Monitor

The following procedure describes how to view job details in the Activity Monitor in the **NetBackup Administration Console**.

To view job details in the Activity Monitor

- ◆ In the **NetBackup Administration Console**, click **Activity Monitor**.
Double-click on a job that is displayed in the **Jobs** tab pane.

Deleting completed jobs in the Activity Monitor

The following procedure describes how to delete a completed job.

To delete completed jobs in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job(s) you want to delete.
- 3 Select **Edit > Delete**.

When a parent job is deleted, all the children jobs are deleted as well. When a child job is deleted that has no children, that child job alone is deleted.

Canceling a job that has not completed in the Activity Monitor

The following procedure describes how to cancel a job that has not completed.

To cancel a job that has not completed in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job that has not completed that you want to cancel. It may be a job that is in the Queued, Re-Queued, Active, Incomplete, or Suspended state.
- 3 Select **Actions > Cancel Job**.

If the selected job is a parent job, all the children of that parent job are canceled as well.

In most cases, a canceled child job cancels only that job and allows the other child jobs to continue. One exception is multiple copies created as part of a policy or a storage lifecycle policy: canceling a child job cancels the parent job and all child jobs.

- 4 To cancel all jobs in the jobs list that have not completed, click **Actions > Cancel All Jobs**.

Restarting a failed (completed) job in the Activity Monitor

The following procedure describes how to restart a job that has completed. Use this procedure to retry a job that has failed.

To restart a completed job in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job that you want to restart.
- 3 Select **Actions > Restart Job**. In this case, a new job ID is created for the job. The job details for the original job reference the job ID of the new job.

Suspending and resuming jobs in the Activity Monitor

The following procedure describes how to suspend restore or backup jobs.

To suspend a restore or a backup job in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job you want to suspend.

Only the backup and the restore jobs that contain checkpoints can be suspended.
- 3 Select **Actions > Suspend Job**.

The following procedure describes how to resume suspended or incomplete jobs.

To resume a suspended or an incomplete job in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor and select the **Jobs** tab.
- 2 Select the suspended or the incomplete job you want to resume.

Only the backup and the restore jobs that contain checkpoints can be suspended.
- 3 Select **Actions > Resume Job**.

Changing the Job Priority dynamically from the Activity Monitor

To dynamically change the priority of a job, select one or more queued or active jobs that wait for resources. Then, either from the **Actions** menu or by right-clicking the job, select **Change Job Priority**.

Select one of the following methods to change the job priority.

Table 28-2 Change Job Priority options

Option	Description
Set Job Priority to	Enters the specific job priority for the selected jobs.
Increment the Job Priority by	Raises the priority of the job by the selected internal.
Decrement the Job Priority by	Lowers the priority of the job by the selected internal.

Changes in the **Change job priority** dialog box affect the priority for the selected job only, and not all other jobs of that type.

To change the job priority defaults, use the **Default Job Priorities** host properties.

See [“Default job priorities properties”](#) on page 85.

About the Daemons tab

The **Daemons** tab in the **Activity monitor** displays the status of NetBackup daemons on primary and media servers.

[Table 28-3](#) describes the NetBackup daemons.

Table 28-3 NetBackup daemons

Daemon	Description
NetBackup Agent Request Server (nbars or nbars.exe)	Populates the NetBackup catalog database with database agent metadata and service requests for agents. This service is also responsible for initiating certain actions, such as starting jobs for Oracle cloning.
NetBackup Audit Manager (nbaudit or nbaudit.exe)	The Audit Manager provides the mechanism to query and report on auditing information.
NetBackup Authentication (nbatd or nbatd.exe)	NetBackup Product Authentication validates identities and forms the basis for authorization and access control in NetBackup. The authentication service also generates security certificates consumed by various NetBackup components. See “About security certificates for NetBackup hosts” on page 36.
NetBackup Authorization (nbazd or nbazd.exe)	NetBackup Product Authorization provides access control in NetBackup applications.
NetBackup Bare Metal Restore Boot Server Service (bmrbd or bmrbd.exe)	Is present if Bare Metal Restore Boot Server is installed. BMR boot servers manage and provide the resources that are used to rebuild systems.
NetBackup Bare Metal Restore Master Server (bmrdr or bmrdr.exe)	Is present if Bare Metal Restore is installed. Manages restoration data, objects, and servers.
NetBackup BMR MTFTP Services (PXEMTFTP or PXEMTFTP.exe)	Is present if Bare Metal Restore is installed. Provides the TFTP protocol services to Bare Metal Restore clients.
NetBackup BMR PXE Service (bmrpxeserver or bmrpxeserver.exe)	Is present if Bare Metal Restore is installed. Provides the PXE protocol services to Bare Metal Restore clients.

Table 28-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Client Service (<i>bpcd</i> or <i>bpcd.exe</i>)	The NetBackup Client daemon. This process issues requests to and from the primary server and the media server to start programs on remote hosts. On UNIX clients, <i>bpcd</i> can only be run in standalone mode. On Windows, <i>bpcd</i> always runs under the supervision of <i>bpinetd.exe</i>. NetBackup has a specific configuration parameter for <i>bpcd</i>. If the port number is changed within the NetBackup configuration, the software also updates the port number in the services file.
NetBackup Compatibility Service (<i>bpcompatd</i> or <i>bpcompatd.exe</i>)	Communicates with legacy NetBackup services.
NetBackup Database Manager (<i>bpdbm</i> or <i>bpdbm.exe</i>)	Manages the NetBackup internal databases and catalogs. <i>BPDBM</i> must be running on the NetBackup primary server during all normal NetBackup operations.
NetBackup Deduplication Engine (<i>spoold</i> or <i>spoold.exe</i>)	Runs on the NetBackup deduplication storage server host to store and manage deduplicated client data. <i>spoold</i> stands for storage pool daemon; do not confuse it with a print spooler daemon. Active only if the NetBackup Data Protection Optimization Option is licensed and the media server is configured as a deduplication storage server.
NetBackup Deduplication Manager (<i>spad</i> or <i>spad.exe</i>)	Runs on the NetBackup deduplication storage server host to maintain the NetBackup deduplication configuration, control deduplication internal processes, control replication, control security, and control event escalation. Active only if the NetBackup Data Protection Optimization Option is licensed and the media server is configured as a deduplication storage server.
NetBackup Device Manager (<i>ltid</i> or <i>ltid.exe</i>)	Starts the Volume Manager (<i>vmd</i>), the automatic volume recognition process (<i>avrd</i>), and any robotic processes. Processes the requests to mount and dismount tapes in robotically controlled devices through the robotic control processes. Mounts the volumes on the tape devices in response to user requests.
NetBackup Enterprise Media Manager (<i>nbemm</i> or <i>nbemm.exe</i>)	Accesses and manages the database where media and device configuration information is stored (<i>EMM_DATA.db</i>). <i>nbemm.exe</i> must be running in order for jobs to run.
NetBackup Event Management Service (<i>nbevtmgr</i> or <i>nbevtmgr.exe</i>)	Provides the communication infrastructure to pass information and events between distributed NetBackup components. Runs on the same system as the NetBackup Enterprise Media Manager.
NetBackup Indexing Manager (<i>nbim</i> or <i>nbim.exe</i>)	Manages the Hold Service.

Table 28-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Job Manager (<code>nbgm</code> or <code>nbgm.exe</code>)	Accepts the jobs that the Policy Execution Manager (<code>nbpem</code> or <code>nbpem.exe</code>) submits and acquires the necessary resources. The Job Manager then starts the job and informs <code>nbpem</code> that the job is completed.
NetBackup KMS (<code>nbkms</code> or <code>nbkms.exe</code>)	A primary server-based symmetric Key Management Service that provides encryption keys to media server BPTM processes.
NetBackup Legacy Client Service (<code>bpnetd</code> or <code>bpnetd.exe</code>)	Listens for connections from NetBackup servers in the network and when an authorized connection is made, starts the necessary NetBackup process to service the connection. Note: On Windows, the Client Service must be run as either an Administrator or Local System account. Problems arise if the Client Service logon account differs from the user that is logged on to use NetBackup. When NetBackup tries to contact the Client Service, a message appears that states the service did not start because of improper logon information. The event is recorded in the Windows System event log. The log notes that the account name is invalid, does not exist, or that the password is invalid. To configure a BasicDisk storage unit that uses CIFS, <code>nbrmms</code> must share the same logon credentials as <code>bpnetd</code> on the media server. See “Configure the NetBackup service credentials for CIFS storage and disk storage units” on page 556.
NetBackup Policy Execution Manager (<code>nbpem</code> or <code>nbpem.exe</code>)	Creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, <code>nbpem</code> is notified and the Policy/Client task objects are updated.
NetBackup Proxy Service (<code>nhostpxy</code> or <code>nhostpxy.exe</code>)	Executes the OpenStorage (OST) calls on any host and returns the results to the proxy plug-in side. The plug-in side returns them to the application. The proxy server (and plug-in) provides a network connection between different servers to relay OpenStorage calls.
NetBackup Remote Manager and Monitor Service (<code>nbrmms</code> or <code>nbrmms.exe</code>)	Discovers and monitors disk storage on NetBackup media servers. Also discovers, monitors, and manages Fibre Transport (FT) connections on media servers and clients for the NetBackup SAN Client option. Runs on NetBackup media servers. To configure a BasicDisk storage unit that uses CIFS, <code>nbrmms</code> must share the same logon credentials as <code>bpnetd</code> on the media server. See “Configure the NetBackup service credentials for CIFS storage and disk storage units” on page 556.

Table 28-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Remote Network Transport Service (<i>nbrntd</i> or <i>nbrntd.exe</i>)	Manages the socket connections between a NetBackup media server and a client that is configured for resilient communication. This service runs on the NetBackup primary server, NetBackup media servers, and clients. NetBackup starts this service when resilient connections are required between hosts. The service stops when resilient connections are no longer required. One instance of the service can process 256 connections. Multiple instances of the service can run simultaneously. See “NBRNTD_IDLE_TIMEOUT option for NetBackup servers” on page 264. See “Resilient network properties” on page 145.
NetBackup Request Daemon (<i>bprd</i> or <i>bprd.exe</i>)	Processes the requests from NetBackup clients and servers. <i>bprd</i> also prompts NetBackup to perform automatically scheduled backups. <i>bprd</i> must be running on the NetBackup primary server to perform any backups or restores.
NetBackup Resource Broker (<i>nbrb</i> or <i>nbrb.exe</i>)	Allocates the storage units, tape drives, and client reservations for jobs. <i>nbrb</i> works with the Enterprise Media Manager (NBEMM). The <i>nbrbutil</i> utility can be used to add or change the Resource Broker settings. See “Using the nbrbutil utility to configure the NetBackup Resource Broker” on page 1017.
NetBackup SAN Client Fibre Transport Service (<i>nbftclnt.exe</i>)	Runs on NetBackup SAN clients. Implements the client side of the Fibre Transport (FT) mechanism. The client FT service opens and closes FT connections and manages the FT connections for shared memory data transfers.
NetBackup Relational Scale-Out Database Manager (<i>pg_ctl.exe</i>)	Manages the NetBackup relational database. The service must be running on the NetBackup primary server during all normal NetBackup operations. On Windows, the service name is <i>vrtsdbsvc_psql</i>.
NetBackup Scale-Out Relational Database Connection Pool Service	Manages the NetBackup scale-out relational database connection pooler provided by PgBouncer. This service must be running on the NetBackup primary server during all normal NetBackup operations.
NetBackup Service Layer (<i>nbsl</i> or <i>nbsl.exe</i>)	Facilitates the communication between the NetBackup graphical user interface and NetBackup logic. This service also does the following for Cloud storage: ■ Generates the metering information for the metering plug-in. ■ Controls the network bandwidth usage with the help of the throttling plug-in. Note: For media servers, the metering service and the throttling service functions are handled by the NetBackup CloudStore Service Container (<i>nbcssc</i>) service.

Table 28-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Service Monitor (nbsvcmon or nbsvcmon.exe)	Monitors the NetBackup services that run on the local computer. If a service unexpectedly terminates, the service tries to restart the terminated service. If <code>nbsvcmon</code> determines that NetBackup is configured for a cluster, the service shuts down, and the monitoring is taken over by the cluster.
NetBackup Storage Lifecycle Manager (nbstserv or nbstserv.exe)	Manages the storage lifecycle operations and schedules duplication jobs. Monitors the disk capacity on capacity-managed volumes and removes older images when required. The SLP Manager and the Import Manager run within <code>nbstserv</code>: ■ The SLP Manager creates batches of the images to be imported based on SLP name and storage device (disk media ID or robot number). ■ The Import Manager monitors a worklist in EMM for images to be imported and initiates <code>bpimport</code> jobs for those images. Note: Restart <code>nbstserv</code> after making changes to the underlying storage for any operation in an SLP.
NetBackup Vault Manager (nbvault or nbvault.exe)	Manages NetBackup Vault. <code>NBVAULT</code> must be running on the NetBackup Vault server during all NetBackup Vault operations.
NetBackup Volume Manager (vmd or vmd.exe)	Manages the volumes (tapes) needed for backup or restore and starts local device management daemons and processes.
NetBackup Web Management Console (nbwmc or nbwmc.exe)	The process for the NetBackup Web Management Console. Manages requests for certificate and host management, and Cloud storage configuration. To configure ports for the NetBackup Web Services, see "Configuring ports for the NetBackup Web Services" in the NetBackup Network Ports Reference Guide.
Veritas Private Branch Exchange (pbx_exchange.exe)	Provides single-port access to clients outside the firewall that connect to NetBackup services. Service name: <code>VRTSpxb</code>. Note: This service does not appear in the Activity Monitor but is represented in the Windows Services utility.

Using the nbrbutil utility to configure the NetBackup Resource Broker

The NetBackup Resource Broker (`nbrb`) allocates resources and maintains resource requests for jobs in the job queue. Use the `nbrbutil` utility to configure the Resource Broker.

The `nbrbutil` utility is located in the following directory:

- On Windows: `install_path\NetBackup\bin\admincmd\nbrbutil`

- On UNIX: `/usr/opensv/netbackup/bin/admincmd/nbrbutil`

For a complete description of `nbrbutil`, see the [NetBackup Commands Reference Guide](#).

[Table 28-4](#) describes the options available to `nbrbutil` command.

Table 28-4 `nbrbutil` options

Option	Description
<code>-cancel requestID</code>	Cancels the allocation request within the given identifier.
<code>-changePriority requestID</code>	Changes the request priority.
<code>-changePriorityClass requestID</code> <code>-priorityClass priorityClass</code>	Changes the request priority class.
<code>-changeSettings</code> <code>parameterparameter_value</code>	Adds or changes the <code>nbrb</code> configuration settings. Table 28-5 describes the configuration settings in detail.
<code>-deleteSetting settingname</code>	Deletes a Resource Broker configuration setting that is identified by <code>settingname</code> .
<code>-dump</code>	Dumps all Resource Broker allocation and request lists.
<code>-dumptables [-f filename]</code>	Enables the Resource Broker to log its internal state in the specified file name.
<code>-help</code>	Lists the help for this command.
<code>-listActiveDriveJobs</code> <code>[driveName]</code>	Lists all the active jobs for a drive.
<code>-listActiveJobs</code>	Lists all the active jobs.
<code>-listActiveMediaJobs mediaId</code>	Lists all the active jobs for a media ID (disk or tape).
<code>-listActivePoolJobs poolName</code>	Lists all the active jobs for a volume pool.
<code>-listActiveStuJobs</code> <code>stuName stugroup</code>	Lists all the active jobs for a storage unit or a storage unit group.
<code>-listOrphanedDrives</code>	Lists the drives that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-listOrphanedMedia</code>	Lists the media that is reserved in EMM but has no corresponding allocation in the Resource Broker.
<code>-listOrphanedPipes</code>	Lists the orphaned Fibre Transport pipes.

Table 28-4 nbrbutil options (*continued*)

Option	Description
<code>-listOrphanedStus</code>	Lists the storage units that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-listSettings</code>	Lists the configuration settings of the Resource Broker.
<code>-priority <i>priority</i></code>	Changes the request priority.
<code>-release <i>allocationID</i></code>	Release the allocation with the given identifier.
<code>-releaseAllocHolds</code>	Releases the allocation holds caused by allocation errors for drives and media.
<code>-releaseDrive <i>drivename</i></code>	Releases all allocations for the named drive.
<code>-releaseMDS <i>mdsAllocationKey</i></code>	Releases the EMM and the MDS allocations that MDS allocates by the specified identifier.
<code>-releaseMedia <i>mediaid</i></code>	Releases all allocations for the specified volume.
<code>-releaseOrphanedDrive <i>drivekey</i></code>	Releases the drives that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-releaseOrphanedMedia <i>mediakey</i></code>	Releases the media that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-releaseOrphanedPipes</code>	Releases the orphaned Fibre Transport pipes.
<code>-releaseOrphanedStu <i>stuName</i></code>	Releases the storage units that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-reportInconsistentAllocations</code>	Reports inconsistent the allocations between the Resource Broker and MDS.
<code>-resetAll</code>	Resets all Resource Broker allocations, requests, and persisted states.
<code>-resetMediaServer <i>mediaserver</i></code>	Resets all Resource Broker EMM and MDS allocations that are related to <i>toItid</i> on the media server.
<code>-resume</code>	Resumes the Resource Broker processing.
<code>-setDriveGroupUnjoinable</code>	Disables the future job from joining the group for this drive.
<code>-setMediaGroupUnjoinable</code>	Disables the future job from joining the group for this media.
<code>-suspend</code>	Suspends the Resource Broker processing.
<code>-syncAllocations</code>	Syncs up any allocation difference between the Resource Broker and MDS.

Table 28-5 lists the parameters for the `nbrbutil -changesettings` option, and describes the use of each.

Use the `nbrbutil` command with the `-changesettings` option to add or change Resource Broker configuration settings.

Table 28-5 `nbrbutil -changesettings` parameters

Parameter	Description
<code>RB_DO_INTERMITTENT_UNLOADS</code>	When the <code>RB_DO_INTERMITTENT_UNLOADS</code> parameter is set to <i>true</i> (default), <code>nbrb</code> initiates unloads of the drives that have exceeded the media unload delay. Drives become available more quickly to jobs that require different media servers or different media than the job that last used the drive. However, the loaded media or drive pair may not be available for jobs with less priority in the prioritized evaluation queue that can use the drive or media without unload. <code>RB_DO_INTERMITTENT_UNLOADS=true</code>
<code>RB_ENABLE_OPTIMIZATION</code>	When the <code>RB_ENABLE_OPTIMIZATION</code> parameter is set to <i>true</i> (default), this entry instructs <code>nbrb</code> to cache states of resource requests. <code>RB_ENABLE_OPTIMIZATION=true</code>
<code>RB_RESPECT_REQUEST_PRIORITY</code>	When the <code>RB_RESPECT_REQUEST_PRIORITY</code> parameter is set to <i>false</i> (default), <code>nbrb</code> continues to evaluate jobs in the prioritized job queue. As a result, a job is likely to reuse a drive more quickly after the drive has been released. However, some lower priority jobs may receive drives before higher priority jobs do. When the <code>RB_RESPECT_REQUEST_PRIORITY</code> parameter is set to <i>true</i>, <code>nbrb</code> restarts its evaluation queue at the top of the prioritized job queue after resources have been released. <code>RB_RESPECT_REQUEST_PRIORITY=false</code>
<code>RB_BREAK_EVAL_ON_DEMAND</code>	When a high priority request appears (a tape span request, a subsequent request for a synthetic or a duplication job, or a read request for an optimized duplication), <code>nbrb</code> immediately interrupts the evaluation cycle. <code>nbrb</code> releases and unloads drives, if required, before a new evaluation cycle is started. If the <code>RB_BREAK_EVAL_ON_DEMAND</code> parameter is set to <i>true</i> (default), the evaluation cycle can be interrupted by high priority requests. <code>RB_BREAK_EVAL_ON_DEMAND=true</code>

Table 28-5 `nbrbutil -changesettings` parameters (*continued*)

Parameter	Description
<code>RB_MAX_HIGH_PRIORITY_QUEUE_SIZE</code>	Spanning requests and additional resources for an active duplication job are put in a special queue for priority processing. The <code>RB_MAX_HIGH_PRIORITY_QUEUE_SIZE</code> parameter sets the maximum number of requests that NetBackup allows in that queue. (Default: 100 requests.) <code>RB_MAX_HIGH_PRIORITY_QUEUE_SIZE=100</code>
<code>RB_RELEASE_PERIOD</code>	The <code>RB_RELEASE_PERIOD</code> parameter indicates the interval that NetBackup waits before it releases a resource. (Default: 180 seconds.) <code>RB_RELEASE_PERIOD=180</code>
<code>RB_CLEANUP_OBSOLETE_DBINFO</code>	The <code>RB_CLEANUP_OBSOLETE_DBINFO</code> parameter indicates the number of seconds that can elapse between the cleanup of obsolete information in the <code>nbrb</code> database. (Default: 60 seconds.) <code>RB_CLEANUP_OBSOLETE_DBINFO=60</code>
<code>RB_MPX_GROUP_UNLOAD_DELAY</code>	The <code>RB_MPX_GROUP_UNLOAD_DELAY</code> parameter indicates the number of seconds that <code>nbrb</code> waits for a new job to appear before a tape is unloaded. (Default: 10 seconds.) <code>RB_MPX_GROUP_UNLOAD_DELAY=10</code> This setting can help avoid unnecessary reloading of tapes and applies to all backup jobs. During user backups, <code>nbrb</code> uses the maximum value of <code>RB_MPX_GROUP_UNLOAD_DELAY</code> and the Media unmount delay host property setting when <code>nbrb</code> unmounts the tape. During restores, Media unmount delay is used, not <code>RB_MPX_GROUP_UNLOAD_DELAY</code> . See "Timeouts properties" on page 175.
<code>RB_RETRY_DELAY_AFTER_EMM_ERR</code>	The <code>RB_RETRY_DELAY_AFTER_EMM_ERR</code> parameter indicates how long NetBackup waits after an EMM error before it tries again. The error must be one where a retry is possible. For example, if a media server is down. (Default: 60 seconds.) <code>RB_RETRY_DELAY_AFTER_EMM_ERR=60</code>
<code>RB_REEVAL_PENDING</code>	The <code>RB_REEVAL_PENDING</code> parameter indicates the number of seconds that can elapse between evaluations of the pending request queue. For example, a pending request queue can include, jobs awaiting resources. (Default: 60 seconds.) <code>RB_REEVAL_PENDING=60</code>

Table 28-5 nbrbutil -changesettings parameters (continued)

Parameter	Description
RB_REEVAL_PERIOD	The RB_REEVAL_PERIOD parameter indicates the time between evaluations if an outstanding request is not satisfied, and if no other requests or resources have been released. (Default: Five minutes must pass before the initial request is reevaluated.) RB_REEVAL_PERIOD=300

Types of NetBackup daemons

The following table describes additional information about NetBackup daemons, found on the UNIX platform.

Standalone daemons	Always run and listen to accept connections. Examples include <code>bpdbm</code> , <code>bprd</code> , <code>bpjobjd</code> , and <code>vmd</code> .
Multiprocess standalone daemons	Splits or forks a child process to handle requests. Examples include <code>bpdbm</code> and <code>bprd</code> .
Single-process standalone daemons	Accept connections and handle requests in the same process.
<code>inetd</code> daemons	<code>inetd(1m)</code> or <code>bpinetd</code> usually launch these NetBackup daemons. Examples include <code>bpcd</code> , <code>bpjava-msvc</code> , and <code>vnetd</code> .

It is recommended that you exit all instances of the **NetBackup Administration Console** after restarting daemons in the **Activity Monitor** or by using a command. Then restart the console with the `jnbSA` command.

The `jnbSA` command is described in the [NetBackup Commands Reference Guide](#).

Monitoring NetBackup daemons

The following procedure describes how to monitor NetBackup daemons.

To monitor NetBackup daemons

- 1 In the **NetBackup Administration Console**, select **Activity Monitor**.
- 2 Select the **Daemons** tab.

- 3 Double-click the name of the daemon to view the details.
- 4 In the **Daemon Details** dialog box, click the up or down arrow to see the details of the next or the previous daemon in the list.

For a description of the daemon, click **Help** in the dialog box.

Starting or stopping a daemon

The following procedure describes how to start or stop a NetBackup daemon.

To start or stop a NetBackup daemon

- 1 In the **NetBackup Administration Console**, select **Activity Monitor**.
- 2 Select the **Daemons** tab.
- 3 Select the daemon(s) that you want to start or stop.
- 4 Select **Actions > Stop Selected** or **Actions > Start Selected**.

Or, select **Actions > Start Daemon** or **Actions > Stop Daemon**.

To start or stop daemons requires the necessary user permissions on the system where the daemon runs.

Displaying all media servers in the Activity Monitor

The **Activity Monitor** may not immediately display all media servers in the **Daemons** tab as soon as the media is added. Also, the **Media Servers** tab of the **Storage Server** dialog may not immediately display all available media servers in a cloud environment.

Even though the media servers may not be visible, it does not affect existing media servers or other NetBackup operations.

To display all media servers:

- Close the **NetBackup Administration Console** on the primary server.
- Stop and restart the NetBackup Service Layer (NBSL). Restarting NBSL does not affect any ongoing backup or restore jobs.
- Open the **NetBackup Administration Console**. The services of the newly added media servers should be visible in the **NetBackup Administration Console**.

This situation affects only the **Remote Administration Console** on Windows.

About the Processes tab

The **Processes** tab in the **Activity monitor** displays the NetBackup processes that run on primary and media servers.

[Table 28-6](#) lists and describes the NetBackup processes.

Table 28-6 NetBackup processes

Process	Port	Description
acsd	13702	The <code>acsd</code> (Automated Cartridge System) daemon runs on the NetBackup media server and communicates mount and unmount requests to the host that controls the ACS robotics.
acsse1	None	On UNIX: The NetBackup ACS storage server interface (SSI) event logger <code>acsse1</code> logs events.
acsssi	None	On UNIX: The NetBackup ACS storage server interface (SSI) <code>acsssi</code> communicates with the ACS library software host. <code>acsssi</code> processes all RPC communications from <code>acsd</code> or from the ACS robotic test utility that is intended for the ACS library software.
avrd	None	The Automatic Volume Recognition process handles automatic volume recognition and label scans. The process allows NetBackup to read labeled tapes and assign the associated removable media requests to drives.
bmrtd	8362	The process for the NetBackup Bare Metal Restore Master Server service.
bpcd	13782	The NetBackup Client daemon issues requests to and from the primary server and the media server to start programs on remote hosts. On UNIX clients, <code>bpcd</code> can only be run in standalone mode. On Windows, <code>bpcd</code> always runs under the supervision of <code>bpinetd.exe</code>. NetBackup has a specific configuration parameter for <code>bpcd</code>: if the port number is changed within the NetBackup configuration, the software also updates the port number in the services file.
bpcompatd	None	The process for the NetBackup Compatibility service.
bpdbm	13721	The process for the NetBackup Database Manager service. The process that responds to queries that are related to the NetBackup catalog. Manages the NetBackup internal databases and catalogs. This service must be running on the NetBackup primary server during all normal NetBackup operations.

Table 28-6 NetBackup processes (*continued*)

Process	Port	Description
bpinetd	None	On Windows: The process for the NetBackup Legacy Client Service. The process that provides a listening service for connection requests. Note: To configure a BasicDisk storage unit that uses CIFS, the media server and the following processes must have the same logon credentials: <code>bpinetd</code>, <code>nbrmms</code>, and <code>vnetd</code>. See “Configure the NetBackup service credentials for CIFS storage and disk storage units” on page 556.
bpjava-msvc	None	The NetBackup Java application server authentication service program. <code>bpinetd</code> starts the program during startup of the NetBackup Java applications and authenticates the user that started the NetBackup Java application.
bpjava-susvc	None	The NetBackup Java application server user service program on NetBackup servers. <code>bpjava-msvc</code> starts the program upon successful login with the NetBackup login dialog box. <code>bpjava-susvc</code> services all requests from the NetBackup Java applications for administration and end user operations on the host on which the NetBackup Java application server is running.
bpjobd	13723	The NetBackup Jobs Database Management daemon. This process queries and updates the jobs database.
bprd	13720	The process for the NetBackup Request Daemon. The process that starts the automatic backup of clients and responds to client requests for file restores and user backups and archives. NetBackup has a specific configuration parameter for <code>bprd</code>: if the port number changes within the NetBackup configuration, the software also updates the port number in the services file.
ltid	None	The process for the NetBackup Device Manager service.
nbatd	13783	The NetBackup Authentication Service validates, identifies, and forms the basis for authorization and access.
nbaudit	None	The NetBackup Audit Manager runs on the primary server. The Enterprise Media Manager (EMM) maintains audit records in the NetBackup database. The act of starting or stopping <code>nbaudit</code> is audited, even if auditing is disabled.
nbazd	13722	The NetBackup Authorization Service verifies that an identity has permission to perform a specific task.
nbars	None	The NetBackup Agent Request Server service populates the NetBackup catalog database with database agent metadata and services request for agents. This service is also responsible for initiating certain actions, such as starting jobs for Oracle cloning.

Table 28-6 NetBackup processes (*continued*)

Process	Port	Description
nbemm	None	The process for the NetBackup Enterprise Media Manager service. The process that accesses and manages the database where media and device configuration information is stored. <code>nbemm.exe</code> must be running in order for jobs to run.
nbEvtMgr	None	The process for the NetBackup Event Manager service. The process that creates and manages event channels and objects for communication among NetBackup daemon. The Event Manager daemon runs with the Enterprise Media Manager (<code>nbemm</code>) only on primary servers.
nbfdrv64	None	The process that controls the Fibre Transport target mode drivers on the media server. <code>nbfdrv64</code> runs on the media servers that are configured for NetBackup Fibre Transport.
nbftsrvr	None	The Fibre Transport (FT) server process that runs on the media servers that are configured for NetBackup Fibre Transport. It does the following for the server side of the FT connection: controls data flow, processes SCSI commands, manages data buffers, and manages the target mode driver for the host bus adaptors.
nbjm	None	The process for the NetBackup Job Manager service. The process that accepts the jobs that the Policy Execution Manager (<code>NBPEM</code>) submits and acquires the necessary resources. The Job Manager then starts the job and informs <code>nbpem</code> that the job is completed.
nbpem	None	The process for the NetBackup Policy Execution Manager service. It creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, <code>NBPEM</code> is notified and the appropriate Policy/Client tasks are updated.
nbproxy	None	The process that safely allows multithreaded NetBackup processes to use existing multithreaded unsafe libraries.
nbrb	None	This process allocates storage units, tape drives, and client reservations for jobs. <code>nbrb</code> works with the Enterprise Media Manager (<code>NBEMM</code>).
nbrmms	None	The process for the NetBackup Remote Manager and Monitor service. It enables NetBackup to remotely manage and monitor resources on a system that are used for backup (or affected by backup activity). Note: To configure a BasicDisk storage unit that uses CIFS, the media server and the following processes must have the same logon credentials: <code>bpineta</code> , <code>nbrmms</code> , and <code>vnetd</code> . See "Configure the NetBackup service credentials for CIFS storage and disk storage units" on page 556.

Table 28-6 NetBackup processes (*continued*)

Process	Port	Description
nbsl	9284	The process for the NetBackup Service Layer service. nbsl listens on this port for connections from local processes and then facilitates the communication between the graphical user interface and NetBackup logic. The port was formerly used by visd.
nbstserv	None	The process for the NetBackup Storage Lifecycle Manager. Manages the storage lifecycle policy operations and schedules duplication jobs. Monitors the disk capacity on the volumes that are capacity-managed and removes older images when required. Note: Restart nbstserv after making changes to the underlying storage for any operation in an SLP.
nbsvcmon	None	The process for the NetBackup Service Monitor. Monitors the NetBackup services. When a service unexpectedly terminates, nbsvcmon attempts to restart the terminated service.
nbvault	None	If Vault is installed, the process for the NetBackup Vault Manager service.
nbwmc	None	The process for the NetBackup Web Management Console. Manages requests for certificate and host management, and Cloud storage configuration. Note: For NetBackup release versions up to 8.1.2, the NetBackup CloudStore Service Container (nbcssc) service is used for the cloud storage configuration. To configure ports for the NetBackup Web Services, see "Configuring ports for the NetBackup Web Services" in the NetBackup Network Ports Reference Guide .
ndmp	10000	NDMP is the acronym for Network Data Management Protocol. NDMP servers are designed to adhere to this protocol and listen on port 10000 for NDMP clients to connect to them.
oprdr	None	The NetBackup Volume Manager (vmd) starts the oprdr operator request daemon. This process receives requests to mount and unmount volumes and communicates the requests to the NetBackup Device Manager ltid. The NetBackup Device Manager communicates the requests to the robotics through SCSI interfaces.
pgbouncer	13787	Manages the NetBackup scale-out relational database connection pooler provided by PgBouncer. This service must be running on the NetBackup primary server during all normal NetBackup Backup operations. The Windows service name is vrtspgbouncersvc.

Table 28-6 NetBackup processes (*continued*)

Process	Port	Description
postgres	13785	The NetBackup relational database process. Some users may notice that a number of instances of <code>postgres</code> run on the NetBackup server. One instance is the primary server process that runs the database cluster. It is the first process started and performs recovery operations, initializes shared memory, and runs background processes. PostgreSQL also spawns additional processes when there is a connection request from a client process. (For more information, see the PostgreSQL documentation: https://www.postgresql.org/docs/current/app-postgres.html.) Each background instance is dedicated to specific purpose. For example: automatic database maintenance, logging error messages, updating and collecting statistics, and handling the client connection from the various programs that want to perform database activities.
spad	10102	The NetBackup Deduplication Manager manages the PureDisk Deduplication Engine. Runs on the NetBackup deduplication storage server host to maintain the NetBackup deduplication configuration, control deduplication internal processes, control replication, control security, and control event escalation.
spoold	10082	The process for the NetBackup Deduplication Engine service. It runs on the deduplication storage server. Active only if the NetBackup Data Protection Optimization Option is licensed and configured.
tladd tldcd	13711	The <code>tladd</code> process runs on a NetBackup server that manages a drive in a Tape Library DLT. This process receives NetBackup Device Manager requests to mount and unmount volumes and sends these requests to the robotic-control process <code>tldcd</code>. The <code>tldcd</code> process communicates with the Tape Library DLT robotics through SCSI interfaces. To share the tape library, <code>tldcd</code> runs on the NetBackup server that provides the robotic control.
vmd	13701	The process for the NetBackup Volume Manager service.

Table 28-6 NetBackup processes (*continued*)

Process	Port	Description
vnetd	13724	The process for the Cohesity Network Daemon, which allows all socket communication to take place while it connects to a single port. The following <code>vnetd</code> process and proxy types can exist on NetBackup hosts: ■ Standalone. A standalone process must exist, and more than one can exist. ■ Inbound proxy. An inbound proxy must exist, and more than one can exist, each identified by a different number. ■ Outbound proxy. An outbound proxy must exist, and more than one can exist, each identified by a different number. ■ HTTP tunnel proxy. By default, an HTTP tunnel proxy should run on NetBackup media servers. It does not run on NetBackup clients. See “WEB_SERVER_TUNNEL_ENABLE option for NetBackup servers” on page 328. You can determine the <code>vnetd</code> process and proxy types as follows: ■ On UNIX and Linux, you can use the NetBackup <code>bpps</code> command. ■ On Windows, you can use the Task Manager Processes tab (you must show the Command Line column). Note: To configure a BasicDisk storage unit that uses CIFS, the media server and the following processes must have the same logon credentials: <code>bpinetd</code>, <code>nbrmms</code>, and <code>vnetd</code>. See “Configure the NetBackup service credentials for CIFS storage and disk storage units” on page 556.
veritas_pbx	1556 1557	The Veritas Private Branch Exchange allows all socket communication to take place while it connects through a single port. Connections to NetBackup use the <code>veritas_pbx</code> port.

Monitoring NetBackup processes in the Process Details dialog box

The following procedure describes how to view the details for a process.


To view the details for a process

- 1 In the **NetBackup Administration Console**, click **Activity Monitor**.
- 2 To view the details for a specific process, double-click on the process you want to display in the **Processes** tab. The **Process Details** dialog box appears for the selected process.
- 3 In the **Process Details** dialog box, click the up or down arrow to see the details of the next process in the list.

About the Drives tab

The **Drives** tab in the Activity Monitor displays the status of NetBackup drives on the monitored server. Use the **Drives** tab to manage drives, device paths, and service requests for operators. Change the status of the drive, clean the tape drive, and perform other operations that you can also perform using the **Device Monitor** under **Media and Device Management**.

Figure 28-2 Activity Monitor Drives tab



To view the details for a drive, double-click the drive in the **Drives** tab pane. For a description of the drive details, click **Help** in the **Drives Details** dialog box.

Monitoring tape drives

The following procedure describes how to monitor NetBackup tape drives.

To monitor NetBackup tape drives

- 1 In the **NetBackup Administration Console**, click the **Activity Monitor**.
- 2 In the right pane, select the **Drives** tab. Double-click a drive from the drive list to view a detailed status.
- 3 A **Drives Details** dialog box appears for the drive you selected. To view the status of the previous drive or the next drive, click the up or down arrow.

Cleaning tape drives from the Activity Monitor

Drive cleaning functions can also be performed from the Device Monitor.


To clean a tape drive

- 1 In the **NetBackup Administration Console**, select **Activity Monitor**. Then, select the **Drives** tab in the **Details** pane.
- 2 Select the drive that you want to clean.
- 3 Select **Actions > Drive Cleaning**, then select one of the following drive cleaning actions.

Action	Description
Clean Now	Starts an operator-initiated cleaning of the selected drive, regardless of the cleaning frequency or accumulated mount time. If the drive is a standalone drive, it must contain a cleaning tape for a mount request to be issued. Clean Now resets the mount time to zero, but the cleaning frequency value remains the same.
Reset Mount Time	Resets the mount time for the selected drive to zero. Use Reset Mount Time to reset the mount time after doing a manual cleaning of a drive.
Set Cleaning Frequency	Sets the number of mount hours between drive cleanings.

About the Error Logs tab

The **Error Logs** tab in the **Activity Monitor** displays the status of NetBackup error logs on the monitored server. You can enable viewing the **Error Logs** tab in the Activity Monitor by navigating to **View** and then selecting **Options**. Select the option for enabling the view of error logs. Use the **Error Logs** tab on the **Activity Monitor** to view the critical logs, information logs, errors and warnings generated during the last 'n' hours duration which is configurable in the **Error Logs** tab.

Figure 28-3 Activity Monitor Error Logs tab

Click on a log entry to view the details that are associated with the log entry in a separate **Log details** dialog box. For a description of the drive details, click **Help** in the **Log Details** dialog box.

About the jobs database

NetBackup uses the `bpdjobs -clean` command to periodically delete the jobs that are done.

By default, the `bpdjobs` process deletes all completed jobs that are more than three days old. By default, the `bpdjobs` process retains more recent done jobs until the three-day retention period expires.

To keep jobs in the jobs database longer than the default of three days, you must change the default value.

If the `bprd` NetBackup Request Daemon is active, `bprd` starts the `bpdjobs` process automatically when it performs other cleanup tasks. The process starts the first time `bprd` wakes up after midnight. The automatic startups occur regardless of whether you choose to run `bpdjobs` at other times by using `cron` or alternate methods.

The `bpdbjobs -clean` is located in the following directory:

On Windows: `install_path\NetBackup\bin\admincmd\bpdbjobs -clean`

On UNIX: `/usr/openv/netbackup/bin/admincmd/bpdbjobs -clean`

Changing the default `bpdbjobs_options` values

Use the one of the following methods to change the default values of the `bpdbjobs_options` on a permanent basis.

Note: In the following instructions that an example value of 192 is given, which is equal to 8 days. This value can be replaced with your desired number of hours.

Windows

To add new registry keys to

`HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\Config` perform the following procedure.

To change the default values for `bpdbjobs_options` on Windows

- 1 Navigate to the following directory:

```
install_path\Veritas\NetBackup\bin\admincmd\
```

- 2 To set the number of hours that unsuccessful jobs are kept in the jobs database or Activity monitor display, run:

```
echo KEEP_JOBS_HOURS = 192 | bpsetconfig
```

- 3 To set the number of hours that successful jobs are kept in the jobs database or Activity monitor display, run:

```
echo KEEP_JOBS_SUCCESSFUL_HOURS = 192 | bpsetconfig
```

UNIX

To add new configuration to `/usr/openv/netbackup/bp.conf`, perform the following procedure.

To change the default values for `bpdbjobs_options` on UNIX

- 1 Navigate to the following directory:

```
cd /usr/opensv/netbackup/bin/admincmd/
```

- 2 To set the number of hours that unsuccessful jobs are kept in the jobs database or Activity monitor display, run:

```
echo KEEP_JOBS_HOURS = 192 | bpsetconfig
```

- 3 To set the number of hours that successful jobs are kept in the jobs database or Activity monitor display, run:

```
echo KEEP_JOBS_SUCCESSFUL_HOURS = 192 | bpsetconfig
```

Considerations

Consider the following notes when changing the default values:

- The default value for the options `KEEP_JOBS_SUCCESSFUL_HOURS` and `KEEP_JOBS_HOURS` is 78 hours.
- The retention period values are measured against the time the job ended.
- Information about successful jobs cannot be kept longer than information about unsuccessful jobs. If `KEEP_JOBS_SUCCESSFUL_HOURS` is greater than `KEEP_JOBS_HOURS`, `bpdbjobs` sets `KEEP_JOBS_SUCCESSFUL_HOURS` to equal `KEEP_JOBS_HOURS`.
- If `KEEP_JOBS_SUCCESSFUL_HOURS` is set to 0, `bpjobd` uses the `KEEP_JOBS_HOURS` `bpdbjobs` value instead for successful jobs.
If the `KEEP_JOBS_SUCCESSFUL_HOURS` value is greater than 0 but less than `KEEP_JOBS_HOURS`, `KEEP_JOBS_HOURS` is used for unsuccessful jobs only.

About the `BPDBJOBS_OPTIONS` environment variable

The `BPDBJOBS_OPTIONS` environment variable provides a convenient method to set job retention options with a script. The `bpdbjobs` process determines how long to retain a job by checking for the `BPDBJOBS_OPTIONS` environment variable.

If present on Windows: `BPDBJOBS_OPTIONS` overrides the registry key settings.

If present on UNIX: `BPDBJOBS_OPTIONS` overrides the `bp.conf` settings. To customize the output of `bpdbjobs`, add a `BPDBJOBS_COLDEFS` entry to the `bp.conf` file for each column you want to appear in the output.

The following options can be used to determine the length of time NetBackup retains jobs. The options should be entered in lower case in the `BPDBJOBS_OPTIONS` environmental variable.

Table 28-7 BPDBJOBS_OPTIONS environment variable options

Option	Description
<code>-keep_hours <i>hours</i></code>	Use with the <code>-clean</code> option to specify how many hours <code>bpdbjobs</code> keeps unsuccessfully completed jobs. Default: 78 hours. To keep both successful and both failed jobs longer than the default of 78 hours, <code>keep_successful_hours</code> must be used with <code>keep_hours</code> .
<code>-keep_successful_hours <i>hours</i></code>	Use with the <code>-clean</code> option to specify how many hours <code>bpdbjobs</code> keeps successfully completed jobs. The number of hours must be less than or equal to <code>keep_hours</code> . Values outside the range are ignored. Default: 78 hours.
<code>-keep_days <i>days</i></code>	Use with the <code>-clean</code> option to specify how many days <code>bpdbjobs</code> keeps completed jobs. Default: 3 days.
<code>-keep_successful_days <i>days</i></code>	This value must be less than the <code>-keep_days</code> value. Use with the <code>-clean</code> option to specify how many days <code>bpdbjobs</code> keeps successfully completed jobs. Default: 3 days.

In the following example, a batch file (`cleanjobs.bat`) was used on a Windows server. You can copy the script directly from this document and change as needed.

- The first line specifies how long to keep unsuccessful jobs (24 hours) and successful jobs (five hours).
- The second line specifies the path to the `bpdbjobs` command. Indicate the correct location of `bpdbjobs` in the `.bat` file. In this example, NetBackup was installed in the default location:


```
set BPDBJOBS_OPTIONS= -keep_hours 24 -keep_successful_hours 5  
C:\progra~1\VERITAS\NetBackup\bin\admincmd\bpdbjobs -clean
```

The same script on a UNIX server would look like the following:

```
setenv BPDBJOBS_OPTIONS "-keep_hours 24 -keep_successful_hours 5 -clean"  
/usr/opensv/netbackup/bin/admincmd/bpdbjobs ${*}
```

You can store the `.bat` file anywhere, as long as it is run from the appropriate directory.

In the following example, the administrator created and stored `cleanjobs.bat` in `C:\Program Files\VERITAS\NetBackup`.

Figure 28-4 Running cleanjobs.bat on Windows

```
C:\WINNT\System32\cmd.exe
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1996 Microsoft Corp.

C:\>cd program files\veritas\netbackup

C:\Program Files\VERITAS\NetBackup>cleanjobs.bat

C:\Program Files\VERITAS\NetBackup>set BPDBJOBS_OPTIONS= -keep_hours 24 -keep_successful_hours 5

C:\Program Files\VERITAS\NetBackup>C:\program files\VERITAS\NetBackup\bin\admincmd\bpdbjobs -clean

C:\Program Files\VERITAS\NetBackup>_
```

bpdbjobs command line options

The `bpdbjobs` command interacts with the jobs database to delete or move completed job files. The command line options supersede all other job retention instructions.

The `-clean` option causes **bpdbjobs** to delete the completed jobs that are older than a specified time period as follows:

```
bpdbjobs -clean [ -M <master servers> ]
[ -keep_hours <hours> ] or [ -keep_days <days> ]
[ -keep_successful_hours <hours> ] or
[ -keep_successful_days <days> ]
```

For example, the following command deletes unsuccessful jobs older than 72 hours.

```
bpdbjobs -clean -keep_hours 72
```

More information is available in the [NetBackup Commands Reference Guide](#).

Enabling the bpdbjobs debug log

If you need detailed information on `bpdbjobs` activities, use the following procedure:

Enabling the bpdbjobs debug log

- ◆ Enable the `bpdbjobs` debug log by creating the following directory:

On Windows: `install_path\NetBackup\logs\bpdbjobs`

On UNIX: `/usr/openv/netbackup/logs/bpdbjobs`

Note: Before you use a debug log, read the guidelines about legacy logging in the [NetBackup Logging Reference Guide](#).

About the Device Monitor

Use the **Device monitor** to manage your tape drives, disk pools, and service requests for operators, as follows:

Media mounts	See “About media mount errors” on page 1037.
Pending requests and actions	See “About pending requests and actions” on page 1038.
	See “About pending requests for storage units” on page 1039.
	See “Resubmitting a pending request” on page 1041.
	See “Resolving a pending action” on page 1040.
Tape drives	See “Denying a pending request” on page 1041.
	See “Changing a drive comment” on page 445.
	See “About downed drives” on page 445.
	See “Changing a drive operating mode” on page 446.
	See “Cleaning a tape drive from the Device monitor” on page 450.
	See “Resetting a drive” on page 451.
	See “Resetting the mount time of a drive” on page 451.
	See “Setting the drive cleaning frequency” on page 452.
Disk pools	See “Viewing drive details” on page 452.
	See “Denying a pending request” on page 1041.
	More information about disk pools is available in the NetBackup guide for your disk storage option:
	■ The <i>NetBackup AdvancedDisk Storage Solutions Guide</i>. ■ The <i>NetBackup Cloud Administrator's Guide</i>. ■ The <i>NetBackup Deduplication Guide</i>. ■ The <i>NetBackup OpenStorage Solutions Guide for Disk</i>. ■ The <i>NetBackup Replication Director Solutions Guide</i>.

About media mount errors

Errors can occur when media is mounted for NetBackup jobs. Depending on the type of error, NetBackup adds the mount request to the pending requests queue or cancels the mount request, as follows:

Adds to the pending requests queue	When NetBackup adds the mount request to the queue, NetBackup creates an operator-pending action. The action appears in the Device monitor. A queued mount request leads to one of the following actions: ■ The mount request is suspended until the condition is resolved. ■ The operator denies the request. ■ The media mount time out is reached.
Cancels the request	When a mount request is automatically canceled, NetBackup tries to select other media to use for backups. (Selection applies only in the case of backup requests.) Many conditions lead to a mount request being automatically canceled instead of queued. When a media mount is canceled, NetBackup selects different media so that the backup is not held up.

When NetBackup selects different media

The following conditions can lead to automatic media reselection:

- The requested media is in a DOWN drive.
- The requested media is misplaced.
- The requested media is write protected.
- The requested media is in a drive not accessible to the media server.
- The requested media is in an offline ACS LSM (Automated Cartridge System Library Storage Module). (ACS robot type only.)
- The requested media has an unreadable barcode. (ACS robot type only.)
- The requested media is in an ACS that is not accessible. (ACS robot type only.)
- The requested media is determined to be unmountable.

About pending requests and actions

In the **NetBackup web UI** click **Storage > Device Monitor**. Then click on the **Device monitor** tab. If requests await action or if NetBackup acts on a request, the request displays in the **Pending requests** pane. For example, if a tape mount requires a specific volume, the request displays in the **Pending requests** pane. If NetBackup requires a specific volume for a restore operation, NetBackup loads or requests the volume.

If NetBackup cannot service a media-specific mount request automatically, it changes the request or action to a pending state.

Table 28-8 Pending states

Pending state	Description
Pending request	Specifies that a pending request is for a tape mount that NetBackup cannot service automatically. Operator assistance is required to complete the request. NetBackup displays the request in the Pending requests pane. NetBackup assigns pending status to a mount request when it cannot determine the following: ■ Which standalone drive to use for a job. ■ Which drive in a robot is in Automatic Volume Recognition (AVR) mode.
Pending action	Specifies that a tape mount request becomes a pending action when the mount operation encounters problems, and the tape cannot be mounted. Operator assistance is required to complete the request, and NetBackup displays an action request in the Pending requests pane. Pending actions usually occur with drives in robotic libraries.

About pending requests for storage units

In the **NetBackup web UI**, click **Storage > Device Monitor**. Then click on the **Device monitor** tab.

The following tape mount requests do not appear in the **Pending requests** pane:

- Requests for backups
- Requests for a tape that is required as the target of a duplication operation

These requests are for resources in a storage unit and therefore are not for a specific volume. NetBackup does not assign a mount request for one storage unit to the drives of another storage unit automatically. Also, you cannot reassign the mount request to another storage unit.

If the storage unit is not available, NetBackup tries to select another storage unit that has a working robot. If NetBackup cannot find a storage unit for the job, NetBackup queues the job (a **Queued** state appears in the Activity monitor).

You can configure NetBackup so that storage unit mount requests are displayed in the **Device monitor** if the robot or drive is down. Pending requests display in the **Device monitor**, and you can assign these mount requests to drives manually.

Resolving a pending request

Use the following procedure to resolve a pending request.

To resolve a pending request

- 1 Insert the requested volume in a drive that matches the density of the volume that was requested.
- 2 Open the NetBackup web UI.
- 3 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 4 In the **Pending requests** pane, select the request and note the contents of the following columns of the request:
 - Density
 - Recorded media ID
 - Mode
- 5 Find a drive type that matches the density for the pending request.
- 6 Verify that the drive is up and not assigned to another request.
- 7 Locate the drive. Then ensure that the drive and the pending request are on the same host.
- 8 If necessary, get the media, write-enable it, and insert it into the drive.
- 9 Wait for the drive to become ready, as explained in the vendor's drive equipment manual.
- 10 Locate the request. Then click **Actions > Assign request**.
- 11 Verify that the request was removed from the **Pending requests** pane.
- 12 Click on the drive name, then click on the **Drive status** tab.
Verify that the job request ID appears in the Request ID column for the drive.

Resolving a pending action

A pending action is similar to a pending request. For a pending action, NetBackup determines the cause of the problem and issues an instruction to the operator to resolve the problem.

Use the following procedure to resolve a pending action.

To resolve a pending action

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 3 In the **Pending requests** pane, locate the pending action.

- 4 Click **Actions > Display pending action**.
- 5 Review the list of possible actions and click **OK**.
- 6 Correct the error condition and either resubmit the request or deny the request.
See [“Resubmitting a pending request”](#) on page 1041.
See [“Denying a pending request”](#) on page 1041.

Resubmitting a pending request

After you correct a problem with a pending action, you can resubmit the request.

If the problem is a volume missing from a robot, first locate the volume, insert it into the robot, and then update the volume configuration. Usually, a missing volume was removed from a robot and then requested by NetBackup.

See [“Robot inventory options”](#) on page 539.

To resubmit a request

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 3 In the **Pending requests** pane, locate the request.
- 4 Click **Actions > Resubmit request**.

Denying a pending request

Some situations may require that you deny requests for service. For example, when a drive is not available, you cannot find the volume, or the user is not authorized to use the volume. When you deny a request, NetBackup sends an appropriate status message to the user.

To deny a request

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 3 In the **Pending requests** pane, locate the request.
- 4 Then click **Actions > Deny request**.

Reporting in NetBackup

This chapter includes the following topics:

- [About the Reports utility](#)
- [Running a report](#)
- [Copying report text to another document](#)
- [Saving or exporting a report](#)
- [Printing a report](#)

About the Reports utility

Use the **Reports** utility in the **NetBackup Administration Console** to generate reports to verify, manage, and troubleshoot NetBackup operations. NetBackup reports display information according to job status, client backups, and media contents. Use the **Troubleshooter** to analyze the cause of the errors that appear in a NetBackup report.

In the **Reports** window, in the right pane, manage the report data or select a report to run.

NetBackup offers many different reports to view information about job activity and media.

Table 29-1 NetBackup Reports

Report name	Description
Status of Backups	The Status of Backups report shows status and error information about the jobs that completed within the specified time period. If an error occurred, a short explanation of the error is included in the report.

Table 29-1 NetBackup Reports (*continued*)

Report name	Description
Client Backups	The Client Backups report shows detailed information about the backups that completed within the specified time period.
Problems	The Problems report generates a list of the problems that the server has logged during the specified time period. The information in this report is a subset of the information that is obtained from the All log entries report.
All log entries	The All logs entries report is also available in the NetBackup web UI. The All log entries report generates a list of all log entries for the specified time period. This report includes the information from the Problems report and Media Logs report. This report also displays the transfer rate. The transfer rate is useful to determine rates and predict backup times for future backups. (The transfer rate does not appear for multiplexed backups.)
Images on Media	The Images on Media report generates a list of the media contents as recorded in the NetBackup image catalog. You can generate this report for any type of media (including disk) and filter it according to client, media ID, or path.
Media Logs	The Media Logs report shows the media errors or the informational messages that are recorded in the NetBackup error catalog.
Images on Tape	The Images on Tape report generates the contents of the tape-based media as recorded in the NetBackup image catalog. The Images on Tape is a subset of the Images on Media report.
Tape Logs	The Tape Logs report displays all error logs related to tape-based backup and recovery. The Tape Logs report is a subset of the Media Logs report.
Tape Contents	The Tape Contents report (formerly known as the Media Contents report) generates a list of the contents of a volume as read directly from the media header and backup headers. This report lists the backup IDs (not each individual file) that are on a single volume. If a tape must be mounted, the delay is longer before the report appears. Before running this report, you can choose to override the default job priority for the job. The default priority is specified in the Default Job Priorities host properties.
Tape Summary	The Tape Summary report summarizes active and nonactive volumes for the specified media owner according to expiration date. It also shows how many volumes are at each retention level. In verbose mode, the report shows each media ID and the expiration date. Nonactive media are those with a status of FULL, FROZEN, SUSPENDED, or IMPORTED. Other volumes are considered active. Expired volumes with a status of FULL, SUSPENDED, or IMPORTED do not appear in the report. However, expired volumes with a FROZEN status do appear in the report. NetBackup deletes other expired volumes from the media catalog when it runs backups. An expired volume of a different status can display if the report is run between the time the volume expires and the time that the next backup is done.

Table 29-1 NetBackup Reports (*continued*)

Report name	Description
Tape Written	The Tape Written report identifies the volumes that were used for backups within the specified time period. The report also does not display the volumes that were used for duplication if the original was created before the specified time period.
Tape Lists	The Tape Lists report generates information about the volumes that are allocated for backups for the selected media owner or media ID. This report does not show media for disk type storage units. For the backups that are saved to disk storage units, use the Images on Media report or the Images on Disk report.
Images on Disk	The Images on Disk report generates the image list present on the disk storage units that are connected to the media server. The Images on Disk report is a subset of the Images on Media report, showing only disk-specific columns.
Disk Logs	The Disk Logs report displays all error logs related to disk-based backup and recovery. The Disk Logs report is a subset of the Media Logs report.
Disk Storage Unit Status	The Disk Storage Unit Status report displays the state of the disk storage units in the current NetBackup configuration. (For example, the total capacity and the used capacity of the disk storage unit.) Multiple storage units can point to the same disk pool. When the report query searches by storage unit, the report counts the capacity of disk pool storage multiple times. Storage units that reference disk groups do not display capacity values.
Disk Pool Status	The Disk Pool Status report generates the details of one or more disk pools.

See [“Copying report text to another document”](#) on page 1045.

For information about Vault reports, see the [NetBackup Vault Administrator's Guide](#).

Running a report

The following procedure describes how to run a NetBackup report from the **Reports** utility.

To run a report

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.

NetBackup runs the report for the primary server that is currently selected. To run a report on a different primary server, on the **File** menu, click **Change Server**.

See [“Accessing remote servers”](#) on page 1060.

- 2 In the left pane, click the name of the report you want to run.
For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report. For example:
 - Select the time period that the report should span.
 - Select the media servers and clients on which to run the report.
- 4 Select **Run report**.

Copying report text to another document

The following procedure describes how to copy the text from a NetBackup report and paste it into a spreadsheet or other document.

To copy report text to another document

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.
- 2 In the left pane, double-click the name of the report you want to run.
For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report, and select **Run report**.
- 4 Select the rows of the report you want to copy by holding down the **Shift** or **Ctrl** key.
- 5 On the **Edit** menu, click **Copy to Clipboard**.
- 6 Paste the selected rows into a spreadsheet or other document.

Saving or exporting a report

The following procedure describes how to save or export a NetBackup report.

To save or export a report

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.
- 2 In the left pane, click the name of the report you want to run.

For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report and click **Run Report**.
- 4 On the **File** menu, click **Export**.
- 5 In the **Save As** dialog box, select the location where you want to save the report, and specify the file name.
- 6 Click **Save**.

See [“Copying report text to another document”](#) on page 1045.

Printing a report

The following procedure describes how to print a NetBackup report.

To print a report

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.
- 2 In the left pane, click the name of the report you want to run.

For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report and click **Run Report**.
- 4 On the **File** menu, click **Print**.

Email notifications

This chapter includes the following topics:

- [Send notifications to the backup administrator about failed backups](#)
- [Send notifications to a host administrator about backups](#)
- [Configure the nbmail.cmd script on the Windows hosts](#)
- [Send notifications about KMS certificate expiration](#)

Send notifications to the backup administrator about failed backups

You can send notifications to the backup administrator about backups with a non-zero status.

On UNIX, NetBackup uses the mail transfer agent sendmail to send email notifications. For Windows, NetBackup requires that an application to transfer messages using SMTP is installed and that the `nbmail.cmd` script is configured on the Windows hosts that send notifications.

See [“Configure the nbmail.cmd script on the Windows hosts”](#) on page 1048.

To configure notifications for the backup administrator of a NetBackup host, see the following topic.

See [“Send notifications to a host administrator about backups”](#) on page 1048.

To send notifications to the backup administrator about failed backups

- 1 On the left, select **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary click **Connect**. Then click **Edit primary server**.

- 4 Click **Global attributes**.
- 5 Enter the email address of the administrator. (Separate multiple addresses with commas.)
- 6 Click **Save**.

Send notifications to a host administrator about backups

You can send notifications to the host administrator about successful and failed backups for a specific host.

On UNIX, NetBackup uses the mail transfer agent sendmail to send email notifications. Windows requires that an application to transfer messages with SMTP is installed. You also must configure the `nbmail.cmd` script on the Windows hosts that send notifications.

See [“Configure the nbmail.cmd script on the Windows hosts”](#) on page 1048.

To send notifications for backups of a specific host

- 1 On the left, select **Hosts > Host properties**.
- 2 Select the client.
- 3 If necessary click **Connect**. Then click **Edit client**.
- 4 Click **Universal settings**.
- 5 Choose how to send the email notifications.
 - To send email notifications from the client, select **Client sends email**.
 - To send email notifications from the server, select **Server sends email**.
- 6 Enter the email address of the host administrator. (Separate multiple addresses with commas.)
- 7 Click **Save**.

Configure the nbmail.cmd script on the Windows hosts

For Windows hosts to send and receive email notifications about backups, the `nbmail.cmd` script must be configured on the applicable hosts.

To configure the nbmail.cmd script on the Windows hosts

- 1 Create a backup copy of `nbmail.cmd`.

- 2 On the primary server, locate the following script:

```
install_path\NetBackup\bin\goodies\nbmail.cmd
```

- 3 Copy the script to the following directory on the applicable hosts:

```
install_path\NetBackup\bin\
```

Primary and media server NetBackup sends notifications from the server if you configure the following setting:

- The **Administrator's email address** in Global Attributes.
- The **Server sends email** option in the **Universal Settings**.

Client. NetBackup sends notifications from the client if you configure the following setting:

- The **Client sends email** option in the **Universal Settings**.

- 4 Use a text editor to open `nbmail.cmd`.

The following options are used in the script:

<code>-s</code>	The subject line of the email
<code>-t</code>	Indicates who receives the email.
<code>-i</code>	The originator of the email, though it is not necessarily known to the mail server. The default (<code>-i Netbackup</code>) shows that the email is from NetBackup.
<code>-server</code>	The name of the SMTP server that is configured to accept and relay emails.
<code>-q</code>	Suppresses all output to the screen.

- 5 Adjust the lines as follows:

- Remove **@REM** from each of the five lines to activate the necessary sections for BLAT to run.
- Replace `SERVER_1` with the name of the mail server. For example:

```
@IF "%~4"==" " (
blat %3 -s %2 -t %1 -i Netbackup -server emailserver.company.com -q
) ELSE (
```

```
blat %3 -s %2 -t %1 -i Netbackup -server emailserver.company.com -q -attach %4  
)
```

6 Save nbmail.cmd.

Send notifications about KMS certificate expiration

A notification is generated when the certificate that is used to communicate with the key management service (KMS) server is about to expire.

Following is an example of notification:

The certificate that is used to communicate with the KMS server example,server.com is about to expire in 10 days. If the certificate is not renewed on time, communication with the KMS server fails.

Following is an example of email contents:

Subject: CN=testuser, O=Veritas, OU=safenet

Issuer: CN=InterCA, O=Veritas,OU=safenet

Server: gemalto

Expiry: Wed Sep 21 16:41:32 IST 2022

Days: 973

Administering NetBackup

- [Chapter 31. Management topics](#)
- [Chapter 32. Accessing a remote server](#)
- [Chapter 33. Using the NetBackup Remote Administration Console](#)
- [Chapter 34. Alternate server restores](#)
- [Chapter 35. Managing client backups and restores](#)
- [Chapter 36. Powering down and rebooting NetBackup servers](#)
- [Chapter 37. About Granular Recovery Technology](#)

Management topics

This chapter includes the following topics:

- [Configuring the NetBackup Client Service](#)
- [Units of measure used with NetBackup](#)
- [NetBackup naming conventions](#)
- [Wildcard use in NetBackup](#)

Configuring the NetBackup Client Service

By default, the NetBackup Client Service is configured on Windows with the **Local System** account. The **Local System** account lacks sufficient rights to perform certain backup and restore operations.

For example, for NetBackup to access CIFS volumes, the account must be changed from **Local System** to an account with access to the CIFS share.

To change the NetBackup Client Service logon account on a Windows computer:

- Open the Windows Services application.
- To change the logon account, stop the NetBackup Client Service.
- Open the properties for the NetBackup Client Service.
- Provide the name and password of the account that has the necessary permissions. For example, change the logon to that of *Administrator*.
- Restart the service.

If the logon property is not changed for the NetBackup Client Service, the policy validation fails with status code 4206.

Situations in which the NetBackup Client Service logon account requires changing

The following list contains situations in which the NetBackup Client Service logon account needs to be changed:

- To access CIFS storage for a storage unit.
- To use UNC paths, the network drives must be available to the service account that the NetBackup Client Service logs into at startup. You must change this account on each Windows client that is backed up that contains data that is shared with another computer.
- During a snapshot: To have read access to the share for backup purposes and write access during restores.
The account must be for a domain user that is allowed to access and write to the share. To verify the account, log on as that user and try to access the UNC path. For example: `\\server_name\share_name`.
- For database agents and options, configure the service with a logon account that has the necessary permission or privileges. See the documentation for your agent or option for more information.
- For the database agents that support VMware backups on a NetApp disk array, configure the logon account to one that has access to the disk array.

Units of measure used with NetBackup

For most units of measure for data, NetBackup uses the terms and abbreviations kilobyte (KB), megabyte (MB), and so on to mean the binary, or bitwise, values of each term. NetBackup does not use the powers-of-ten values, such as 1,000 for KB or 1,000,000 for MB.

When you calculate values that appear in NetBackup displays and reports, it is important to understand the difference between a unit's binary value and its powers-of-ten value. For example, a displayed value of 1.5TB actually means 1,649,267,441,664, bytes (the binary value) and not 1,500,000,000,000 bytes (the powers-of-ten value), a difference of almost 150 billion bytes.

The following table shows a number of common displayed units of measure with their corresponding bitwise names, binary multipliers, and actual values.

Table 31-1 Units of measure used in NetBackup

Displayed unit	Bitwise unit	Binary multiplier	Actual value in bytes
Kilobyte (KB)	Kebibyte (KiB)	2 ¹⁰	1024

Table 31-1 Units of measure used in NetBackup (*continued*)

Displayed unit	Bitwise unit	Binary multiplier	Actual value in bytes
Megabyte (MB)	Mebibyte (MiB)	2 ²⁰	1048576
Gigabyte (GB)	Gibibyte (GiB)	2 ³⁰	1073741824
Terabyte (TB)	Tibibyte (TiB)	2 ⁴⁰	1099511627776
Petabyte (PB)	Pebibyte (PiB)	2 ⁵⁰	1125899906842624
Exabyte (EB)	Exbibyte (EiB)	2 ⁶⁰	1152921504606846976

The Institute of Electrical and Electronics Engineers (IEEE) and the International Electrotechnical Commission (IEC) have adopted standards for these values. See the following articles for more information:

- <https://standards.ieee.org/standard/1541-2002.html> (with a paid IEEE subscription)
https://en.wikipedia.org/wiki/IEEE_1541-2002
- https://en.wikipedia.org/wiki/ISO/IEC_80000

NetBackup naming conventions

NetBackup has rules for naming logical constructs, such as clients, disk pools, backup policies, storage lifecycle policies, and so on. Generally, names are case-sensitive. The following set of characters can be used in user-defined names and passwords:

- Alphabetic (A-Z a-z) (names are case-sensitive)
- Numeric (0-9)
- Period (.)
Do not use periods in the WORM volume names.
- Plus (+)
- Hyphen (-)
Do not use a hyphen as the first character.
- Underscore (_)

These characters are also used for foreign languages.

Note: No spaces are allowed.

The Logical Storage Unit (LSU) name or the Domain Volume name must have fewer than 50 ASCII characters including a hyphen (-) and an underscore (_) and must not have a blank space.

Wildcard use in NetBackup

NetBackup recognizes the following wildcard characters in areas where wildcards can be used. (For example, in the paths of include and exclude file lists.)

The following table shows the wildcards that can be used in various NetBackup dialog boxes and lists.

Table 31-2 Wildcard use in NetBackup

Wildcard	Use
*	An asterisk serves as a wildcard for zero or more characters. An asterisk can be used in the backup selection list, the include list, and the exclude list for Windows and UNIX clients. For example: <code>r*</code> refers to all files that begin with <code>r</code> <code>r*.doc</code> refers to all files that begin with <code>r</code> and end with <code>.doc</code>. To back up all files that end in <code>.conf</code>, specify: <code>/etc/*.conf</code>
?	A question mark serves as a wildcard for any single character (A through Z; 0 through 9). A question mark can be used in the backup selection list, the include list, and the exclude list for Windows and UNIX clients. For example: <code>file?</code> refers to <code>file2</code>, <code>file3</code>, <code>file4</code> <code>file??</code> refers to <code>file12</code>, <code>file28</code>, <code>file89</code> To back up all files named <code>log01_03</code>, <code>log02_03</code>, specify: <code>c:\system\log??_03</code>

Table 31-2 Wildcard use in NetBackup (*continued*)

Wildcard	Use
[]	A pair of square brackets indicates any single character or range of characters that are separated with a dash. For example: <code>file[2-4]</code> refers to <code>file2</code>, <code>file3</code>, and <code>file4</code> <code>file[24]</code> refers to <code>file2</code>, <code>file4</code> <code>*[2-4]</code> refers to <code>file2</code>, <code>file3</code>, <code>file4</code>, <code>name2</code>, <code>name3</code>, <code>name4</code> Brackets are not valid wildcards under all circumstances for all clients: ■ Brackets that are used as wildcards in include and exclude lists: Windows clients: Allowed UNIX clients: Allowed ■ Brackets that are used as wildcards in policy backup selections lists: Windows clients: Not allowed; the use of brackets in policy backup selections lists causes backups to fail with a status 71. UNIX clients: Allowed
{ }	Curly brackets can be used in the backup selection list, the include list, and the exclude list for UNIX clients only. A pair of curly brackets (or braces) indicates multiple file name patterns. Separate the patterns by commas only; no spaces are permitted. A match is made for any or all entries. For example: <code>{*1.doc,* .pdf}</code> refers to <code>file1.doc</code>, <code>file1.pdf</code>, <code>file2.pdf</code> Note: Curly brackets are valid characters for Windows file names and cannot be used as wildcards on Windows platforms. Backslashes cannot be used as escape characters for curly bracket characters.

To use wildcard characters literally, precede the character with a backslash (\).

A backslash (\) acts as an escape character only when it precedes a special or a wildcard character. NetBackup normally interprets a backslash literally because a backslash is a legal character to use in paths.

Assume that the brackets in the following examples are to be used literally:

`C:\abc\fun[ny]name`

In the exclude list, precede the brackets with a backslash:

`C:\abc\fun\[ny\]name`

Table 31-3 Placement of wildcards in the path of backup selections

Client type	Examples
For Windows clients, wildcards function correctly only when they are placed at the end of the path, in the file or the directory name. See “Pathname rules for Windows client backups” on page 795.	The following example is allowed: <pre>C:\abc\xyz\r*.doc</pre> Wildcard characters do not work elsewhere in the path. For example, an asterisk functions as a literal character (not as a wildcard) in the following examples: <pre>C:*\xyz\myfile</pre> <pre>C:\abc*\myfile</pre>
For UNIX clients, wildcards can appear anywhere in the path. See “Pathname rules for UNIX client backups” on page 802.	The following examples are allowed: <pre>/etc/*/abc/myfile</pre> <pre>/etc/misc/*/myfile</pre> <pre>/etc/misc/abc/*.*</pre>

See [“Backup Selections tab”](#) on page 782.

Accessing a remote server

This chapter includes the following topics:

- [Prerequisites for accessing a remote server](#)
- [Accessing remote servers](#)
- [Troubleshooting remote server administration](#)

Prerequisites for accessing a remote server

In a NetBackup environment, you can use multiple NetBackup servers to perform various administrative tasks.

For example, consider a NetBackup environment that contains two NetBackup domains. Server1 in one domain and Server2 in the other. You can log into the NetBackup Administration Console of Server1 and then change to Server2 to administer Server2. You must perform the following tasks before you can change to a different server:

- From one server, allow access to another server. See [“Allow access to another server”](#) on page 1058.
- Authorize users of one server to access another server. See [“Authorize users of one server to access another server”](#) on page 1059.

To change to a different server, you can either specify the remote server on the login screen or change to the server from within the NetBackup Administration Console. See [“Accessing remote servers”](#) on page 1060.

Allow access to another server

For a local host to administer a remote server, the name of the local host must appear in the server list of the remote server. Likewise, the remote host must include the local host in its server list.

To add a server to a server list

- 1** In the **NetBackup Administration Console**, expand **Host Properties > Master Server**.
- 2** Double-click the name of the server to view the properties.
- 3** Select the **Servers** tab to display the server list.

If the remote server does not appear in the **Additional Servers** list, the current server considers it invalid.

- 4** To add a server to the server list, click **Add**.
- 5** In the **New Server** dialog box, type the server name in the field.
- 6** Click **Add** to add the server to the list. Then, click **Close** to close the dialog box. The server appears in the server list.

The `bp.conf` file on every UNIX server contains `SERVER` and possibly `MEDIA_SERVER` entries. The server list in the **Servers** properties dialog box represents these entries. Hosts that are listed as media servers have limited administrative privileges.

- 7** Click **OK** to save the changes.

Authorize users of one server to access another server

You must explicitly authorize users of one server to access another server. Add users of a server to the `auth.conf` file on the server that they are authorized to access. In this example, authorize users of Server1 to access Server2. This task must be performed on Server2.

Note: Add only the users that don't have administrative privileges on one server to the `auth.conf` file on another server.

On UNIX, the `auth.conf` file is located at `/usr/opensv/java`.

On Windows, create the `auth.conf` file from the `auth.conf.win.template` file that is located at `C:\Program Files\Veritas\Java`.

To authorize users of one server to access another server

- ◆ In the `auth.conf` file on Server2, add users of Server1 authorized to access Server2.

If Server1 is a Windows host, add the following line in the `auth.conf` file:

For example, `username ADMIN=ALL JBP=ALL`

If Server1 is a UNIX host, add the following line in the `auth.conf` file:

For example, `root ADMIN=ALL JBP=ALL`

Accessing remote servers

If a NetBackup site has multiple primary servers, you can configure the systems so that multiple servers can be accessed from one **NetBackup Administration Console**.

If the server that you want to access is a media server or client, it must be provisioned with a security certificate.

See [“About security certificates for NetBackup hosts”](#) on page 36.

Use the following procedure to access a remote server.

To access a remote server

- 1 Ensure that the remote server is accessible to the local server.
See [“Allow access to another server”](#) on page 1058.
- 2 Use one of the following methods to access a remote server:
 - Use the **Change Server** option:
 - Select any main node in the left pane of the **NetBackup Administration Console**. (The option does not appear in the **File** menu unless a main node is selected.)
 - Select the **File > Change Server** menu command.
 - Click **OK** to launch a new console.

- Enter the host name of another server.
 Enter a user name and password and click **Login**.
 If the user has the necessary permissions on both servers, the user can transition from one to another without setting up trust relationships.
 If *server1* is not listed on the server list of *server2*, *server1* receives an error message after it tries to change servers to *server2*.
 If the user has administrative privileges on one server and different privileges on another server, the user is required to reauthenticate.
- Specify the remote server on the login screen:
 Enter the host name of the remote server.
 Enter the user name and password for an authorized NetBackup administrator, then click **Login**.

To log in to a different server, specify the name of the remote host in the login screen

Troubleshooting remote server administration

To administer a server from another primary server, make sure that the following conditions are met:

- The destination server is operational.
- NetBackup daemons are running on both hosts.
- The network connection is valid.
- The user has administrative privileges on the destination host.
- To perform a **Change Server** operation, all servers should be at the same version.
- The current host is listed in the server list of the destination host.

See [“Allow access to another server”](#) on page 1058.

The host does not need to be listed if the host is a media server or a client. Or, it does not need to be listed if only media and device management or monitoring is to take place.

If you change the primary server list, stop and restart the NetBackup Database Manager (bpdbm) and NetBackup Request Daemon (bprd) to ensure that all appropriate NetBackup processes use the new server entry.

- Authentication is set up correctly, if used.
- To perform a **Change Server** operation to a media server or client, the media server or client must have a security certificate installed.
- For problems changing servers to configure media or devices or monitor devices, verify that the NetBackup Volume Manager is running on that server.
- If you cannot access devices on the remote host, it may be necessary to add a `SERVER` entry to the `vm.conf` file on that host.
See the [NetBackup Administrator's Guide, Volume II](#) for instructions.
- If you cannot start or stop processes or services through the Activity Monitor, verify the following:
 - The remote server is a Windows system. Only on other Windows systems can processes be monitored and controlled.
 - You have the required permissions on the remote server. Windows security must allow access to the user that is running the Activity Monitor.

Using the NetBackup Remote Administration Console

This chapter includes the following topics:

- [About the NetBackup Remote Administration Console](#)
- [Run-time configuration options for the NetBackup Administration Console](#)
- [About improving NetBackup performance](#)
- [About adjusting time zones in the NetBackup Administration console](#)

About the NetBackup Remote Administration Console

The **NetBackup Remote Administration Console** is a stand-alone Java-based administration console. This console is installed on a host that does not have NetBackup server software. It is used to monitor NetBackup servers remotely.

Installing this console installs the **NetBackup Administration Console**. The presence of the client software enables the computer to be backed up like any other client. No primary server software or media server software is installed.

The administrator first starts the **NetBackup Administration Console** using one of the following methods:

- Select **Start > Programs > Veritas NetBackup > NetBackup version Administration Console** on the Windows computer on which the NetBackup Remote Administration Console is installed.

- Run the `jnbSA` command on a UNIX computer where NetBackup is installed.


Then the administrator logs on to the application server on the host that is specified in the logon dialog box.

The application server is the host that is specified in the **NetBackup Administration Console** logon dialog box and authenticates the logon credentials of the user. The credentials are authenticated by using standard UNIX user account data and associated APIs.

Note: The host that is specified in the logon dialog box and the system that runs the **NetBackup Administration Console** must run the same NetBackup version.

Note: To log on to any **NetBackup Administration Console**, your logon credentials must be authenticated from the connecting primary or media server.

Figure 33-1 NetBackup logon dialog box



The server that is usually the object of all administrative tasks is the host that is specified in the **NetBackup Administration Console** logon dialog box.

An exception is the use of the **File > Change Server** capability in the **NetBackup Administration Console**. The **Change Server** capability allows administration of a remote server (a server other than the one specified in the **NetBackup Administration Console** logon dialog box).

Note: To perform a Change Server operation, all servers should be at the same version.

Regardless of which server is administered, all administrative tasks that are performed in the **NetBackup Administration Console** make requests of the application server. All tasks are run on the application server host, whether the server is remote or whether the server is specified on the logon dialog box.

To administrate from a remote server, the application server host must be included in the server list of the remote server.

See [“Allow access to another server”](#) on page 1058.

See [“Accessing remote servers”](#) on page 1060.

Run-time configuration options for the NetBackup Administration Console

On Windows computers, the analogous file containing configuration options for the **NetBackup Administration Console** is `install_path\java\setconf.bat`

On UNIX computers, file `/usr/opensv/java/nbj.conf` contains configuration options for the **NetBackup Administration Console**. Enter one option per line, following the same syntax rules as exist for the `bp.conf` file.

`nbj.conf` and `setconf.bat` contain commands for each of the configuration options that are described in the following topics. To make changes, change the value after the equal sign in the relevant set command.

BROWSER_BINARY_PATH

In some cases, NetBackup may not be able to launch the browser for authentication during the **NetBackup Administration Console** login using the **Single sign-on, Certificates, or Smart Cards through the Web UI** option. If you come across such an error, configure the `BROWSER_BINARY_PATH` option to launch a browser.

This option uses the following format:

```
BROWSER_BINARY_PATH=browser_executable_path
```

For example:

```
BROWSER_BINARY_PATH=/usr/bin/firefox
```

DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT

The `DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT` configuration option specifies the default timeout value of child job for dynamic data streaming for the NAS-Data-Protection policy. After the parent Backup from Snapshot job is triggered, NetBackup starts the NBCS process which pre-processes the child backup jobs. After the pre-processing, NBCS waits for all child jobs to be start after which it allocates a filelist to child jobs for backup. NBCS doesn't start this activity unless all child jobs have started and ready to receive the filelist. By default, the NBCS process waits for 600 seconds for all child jobs to start. Depending on the number of streams per volume set for the NAS-Data-Protection policy and the total number of volumes to be backed up at a given time, the child jobs might take longer to start.

Table 33-1 `DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT` options

Name	<code>DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT</code>
Type	Integer
Default value	600 seconds
Minimum value	300 seconds
Maximum value	3600 seconds

If your scheduled configurations cause such timeout you can change the timeout value by using the configuration parameter `DYNAMIC_STREAMING_START_CHILD_BACKUP_JOBS_TIMEOUT`. You can change the value of this variable by using the `bpsetconfig` command. Use the `bpgetconfig` CLI to view the value of this variable. You can set this configuration parameter on the NetBackup primary server.

FIREWALL_IN

The `FIREWALL_IN` configuration option provides a method to use a **NetBackup Administration Console** that is outside of a trusted network to administer the NetBackup primary servers that are within a trusted network.

This option uses the following format.

On Windows:

```
SET FIREWALL_IN=
HOST1:PORT1=HOST2:PORT2;IP_ADDR1:PORT3=IP_ADDR2:PORT4
SET FIREWALL_IN >> "%NBJDIR%\nbjconf
```

On UNIX:

```
FIREWALL_IN= HOST1:PORT1=HOST2:PORT2[;...;HOSTn:PORTn=HOSTm:PORTm]
```

Where *HOST* is a host name or an IP address.

This configuration option provides a way to allow administrators to bypass the firewall by using one of the following methods:

- Enter the port number of the `bpjava` service in the trusted internal network. Then, map the private interface where the `bpjava` service runs to a public interface that can be reached from outside the firewall.
- Set up a Secure Shell (SSH) tunnel from the local host to the system inside the firewall.

In the following example:

- Primary server `NBPrimary.abc.com` is in a trusted network, behind a firewall.
- The IP address of `NBPrimary.abc.com` is `10.221.12.55`.
- The **NetBackup Administration Console** is installed on localhost.
- SSH tunnels exist from localhost to `NBPrimary.abc.com` as follows:

```
bpjava-msvc port (default 13722)    localhost:port1
vnetd port (default 13724)          localhost:port2
pbx port (default 1556)              localhost:12345
```

Where **localhost** is the host name and `port1` is the IP port.

To make relevant changes for connections to `bpjava-msvc` and `vnetd`, see the following topic:

See [“VNETD_PORT”](#) on page 1073.

On Windows systems, use `setconf.bat` to add the option:

```
SET FIREWALL_IN=
NBMaster.abc.com:1556=localhost:12345;10.221.12.55:12345=localhost:12345
SET FIREWALL_IN >> "%NBJDIR%\nbjconf
```

On UNIX systems, add the following line to the `nbj.conf` file:

```
FIREWALL_IN=NBPrimary.abc.com:1556=localhost:12345;10.221.12.55:12345=localhost:12345
```

The entry indicates the following:

- The connection to NBPrimary.abc.com:1556 is to be redirected to localhost:12345.
- The connection to 10.221.12.55:1556 is to be redirected to localhost:12345.

Note: The same options are used if NBPrimary.abc.com has a public interface (NBPrimarypub.abc.com) that can be reached from the Internet. In this case, the administrator replaces localhost with NBPrimaryPub.abc.com.

FORCE_IPADDR_LOOKUP

The `FORCE_IPADDR_LOOKUP` configuration option specifies whether NetBackup performs an IP address lookup to determine if two host name strings are indeed the same host. This option uses the following format:

```
FORCE_IPADDR_LOOKUP = [ 0 | 1 ]
```

Where:

- 0 Indicates that no IP address lookup is performed to determine if two host name strings are indeed the same host. They are considered to be the same host if the host name strings compare equally. Or, if a short name compares equally to the short name of a partially or fully qualified host name.
- 1 Indicates that an IP address lookup is performed if the two host name strings do not match. The lookup determines if they have the same host. The default is to perform an IP address lookup if necessary to resolve the comparison. The IP address lookup is not performed if the host name strings compare equally.

Note: Use a value of 1 for this option if you have the same host name in two different domains. For example, `eagle.abc.xyz` and `eagle.def.xyz` or by using host name aliases.

Many places in the **NetBackup Administration Console** compare host names to determine if the two are the same host. For example, the **File > Change Server** command.

The IP address lookup can consume time and result in slower response time. However, accurate comparisons are important.

No IP address lookup is necessary if the host name is specified consistently in the **NetBackup Administration Console** logon dialog box. It must match how the host names are configured in NetBackup. Host names are identified in the server list

that is found in the **Servers** host properties. On UNIX systems, the host names also appear in the `bp.conf` file.

Using host names *eagle* and *hawk*, the following describes how this option works:

`FORCE_IPADDR_LOOKUP = 0`

Comparisons of the following result in no IP address lookup. The hosts are considered to be the same host.

```
eagle and eagle
eagle.abc.def and eagle.abc.def
eagle.abc and eagle.abc.def
eagle and eagle.abc.def
eagle and eagle.anything
```

The hosts are considered to be different for any comparisons of short, partially, or fully qualified host names of *eagle* and *hawk* regardless of aliases.

`FORCE_IPADDR_LOOKUP = 1`

Comparisons of the following result in no IP address lookup. The hosts are considered to be the same host.

```
eagle and eagle
eagle.abc and eagle.abc
eagle.abc.def and eagle.abc.def
```

In addition to all comparisons of *eagle* and *hawk*, the following result in an IP address lookup. The comparison determines if the hosts are indeed the same host.

```
eagle.abc and eagle.abc.def
eagle and eagle.abc.def
eagle and eagle.anything
```

INITIAL_MEMORY, MAX_MEMORY

Both `INITIAL_MEMORY` and `MAX_MEMORY` allow configuration of memory usage for the Java Virtual Machine (JVM).

It is recommended that all of the interfaces (the **NetBackup Remote Administration Console**, the **NetBackup Administration Console**, or the **NetBackup Backup, Archive, and Restore** user interface) run on a system that contains at least 1 gigabyte of physical memory. Make sure that 512 megabytes of memory are available to the application.

`INITIAL_MEMORY` specifies how much memory is allocated for the heap when the JVM starts. The value probably does not require changing. The default is sufficient for quickest initialization of `jnbSA`, the **Remote Administration Console**, or `jbpSA` on a system with the recommended amount of memory.

On UNIX systems, the initial memory allocation can also be specified as part of the `jnbSA` or `jbpSA` command. For example:

```
jnbSA -ms 256M
```

Default = 256M (megabytes).

`MAX_MEMORY` specifies the maximum heap size that the JVM uses for dynamically allocated objects and arrays. If the amount of data is large, consider specifying the maximum heap size. For example, a large number of jobs in the Activity Monitor.

On UNIX systems, the maximum memory allocation can also be specified as part of the `jnbSA` or `jbpSA` command. For example:

```
jnbSA -mx 512M
```

Default = 512M (megabytes).

MEM_USE_WARNING

The `MEM_USE_WARNING` configuration option specifies the percent of the memory that is used as compared to `MAX_MEMORY`, at which time a warning dialog box appears to the user. Default = 80%. This option uses the following format:

```
MEM_USE_WARNING=80
```

NB_FIPS_MODE

Use the `NB_FIPS_MODE` option to enable the FIPS mode in your NetBackup domain.

By default, the `NB_FIPS_MODE` option is disabled.

To enable the option, use the following format:

```
NB_FIPS_MODE = ENABLE
```

NBJAVA_CLIENT_PORT_WINDOW

The `NBJAVA_CLIENT_PORT_WINDOW` configuration option specifies the range of non-reserved ports on this computer to use for connecting to the NetBackup application server. It also specifies the range of ports to use to connect to the `bpjjobd` daemon from the Activity Monitor in the **NetBackup Administration Console**.

This option uses the following format:

```
NBJAVA_CLIENT_PORT_WINDOW = n m
```

Where:

n Indicates the first in a range of non-reserved ports that are used for connecting to the `bpjava` processes on the NetBackup application server. It also specifies the range of ports to use to connect to the `bpjobjd` daemon or Windows service from the Activity Monitor of the **Remote Administration Console**.

If *n* is set to 0, the operating system determines the non-reserved port to use (default).

m Indicates the last in a range of non-reserved ports that are used for connecting to the **NetBackup Administration Console** or the **Remote Administration Console**.

If *n* and *m* are set to 0, the operating system determines the non-reserved port to use (default).

The minimum acceptable range for each user is 120. Each additional concurrent user requires an additional 120. For example, the entry for three concurrent users might look as follows:

```
NBJAVA_CLIENT_PORT_WINDOW = 5000 5360
```

If the range is not set wide enough, `jnbSA` exits with an error message that states an invalid value has occurred during initialization.

Note: Performance is reduced with the use of `NBJAVA_CLIENT_PORT_WINDOW`.

NBJAVA_CORBA_DEFAULT_TIMEOUT

The `NBJAVA_CORBA_DEFAULT_TIMEOUT` configuration entry specifies the default timeout that is used for most CORBA operations that the **NetBackup Administration Console** performs.

This option is present by default and uses the following format:

```
NBJAVA_CORBA_DEFAULT_TIMEOUT=60
```

The default is 60 seconds.

NBJAVA_CORBA_LONG_TIMEOUT

The `NBJAVA_CORBA_LONG_TIMEOUT` configuration entry specifies the timeout value that the **NetBackup Administration Console** uses in the following areas:

- Device Configuration Wizard
- Disk Pool Configuration Wizard
- Disk Pool Inventory

This option is present by default and uses the following format:

```
NBJAVA_CORBA_LONG_TIMEOUT=1800
```

The default is 1800 seconds.

NETBACKUP_API_CLIENT_CONNECTION_TIMEOUT

The `NETBACKUP_API_CLIENT_CONNECTION_TIMEOUT` configuration option specifies the default timeout value that the **NetBackup Administration Console** uses while it establishes a connection with the NetBackup web API server.

This option is present by default and uses the following format:

```
NETBACKUP_API_CLIENT_CONNECTION_TIMEOUT=180
```

The default is 180 seconds.

See [“NETBACKUP_API_CLIENT_READ_TIMEOUT”](#) on page 1072.

NETBACKUP_API_CLIENT_READ_TIMEOUT

The `NETBACKUP_API_CLIENT_READ_TIMEOUT` configuration option specifies the default timeout value that the **NetBackup Administration Console** uses when it requests the data from the NetBackup web API server.

This option is present by default and uses the following format:

```
NETBACKUP_API_CLIENT_READ_TIMEOUT=1800
```

The default is 1800 seconds.

See [“NETBACKUP_API_CLIENT_CONNECTION_TIMEOUT”](#) on page 1072.

PBX_PORT

The `PBX_PORT` configuration entry specifies the `pbx` port.

This option is present by default and uses the following format:

```
PBX_PORT=1556
```

USE_URANDOM

Enable the `USE_URANDOM` option to specify `/dev/urandom` as the character device to provide cryptographically secure random output in your NetBackup environment.

The default value of the `USE_URANDOM` option is 0. When the `USE_URANDOM` option is set to default, the character device to be used is based on the value of the `NB_FIPS_MODE` option.

If `NB_FIPS_MODE` is enabled, `dev/random` is used.

If `NB_FIPS_MODE` is disabled, `dev/urandom` is used.

To enable the `USE_URANDOM` option, use the following format:

```
USE_URANDOM = 1
```

If `USE_URANDOM` is set to 2 (or is disabled), the `dev/random` character device is used to provide cryptographically secure random output.

VNETD_PORT

The `VNETD_PORT` is the configured port for the `vnetd` daemon process and is registered with the Internet Assigned Number Authority (IANA).

This option uses the following format:

```
VNETD_PORT=13724
```

It is not recommended to change this port. If changes are necessary, make the change on all NetBackup hosts in the relevant NetBackup cluster.

This option is preserved for backward compatibility when the 7.0.1 **NetBackup Administration Console** is used to communicate with a 7.0 NetBackup server.

For more information, see the [NetBackup Installation Guide](#).

The value must be set in the corresponding `setconf.bat` (Windows) or `nbg.conf` (UNIX) configuration option.

About improving NetBackup performance

The most important factor to consider concerning performance issues for the following interfaces is the platform on which the console runs:

- **Remote Administration Console**
- **NetBackup Administration Console**
- NetBackup **Backup, Archive, and Restore** user interface

Regardless of the platform, you can run the administration console from one of the following locations:

- Run it locally on a desktop host (on supported Windows and UNIX platforms)
- Run it remotely and display it back to a desktop host (from supported UNIX platforms)

To provide the best performance, the recommended method for using these consoles is to run the consoles locally on a desktop host. When the consoles are run locally,

they do not exhibit the font and the display issues that can be present in some remote display-back configurations.

About running the NetBackup Administration Console locally

On Windows platforms, select **Start > Programs > Cohesity NetBackup > NetBackup 8.x Administration Console** to start the console.

On supported UNIX platforms, the console is run locally if `jnbSA` or `jbpSA` is entered on the same host on which the console appears. That is, your display environment variable is set to the host on which the `jnbSA` or `jbpSA` commands were entered.

Improvements in Java technology have made remote X-display back potentially viable on some platforms. However, problems continue with certain controls in the consoles. For example, incorrect box operations, sluggish scrolling, and display problems in tables with many rows. More serious issues have also occurred. Consoles can abort and hang because of a Java Virtual Machine (JVM) failure when run in this mode on some platforms. Therefore, it is not recommended to run the consoles in a remote X-display back configuration.

About running a console locally and administering a remote server

The **NetBackup Administration Console** and the **Backup, Archive, and Restore** user console are distributed applications. Both applications consist of two major and separate system processes that can run on different computers. For example: the **NetBackup Administration Console** on one computer and the console's application server – `bpjava` processes on another computer.

The **NetBackup Administration Console** does not need to run on a NetBackup server host. However, the application server must run on this host in order for you to be able to administer NetBackup.

Although the **NetBackup Administration Console** does not run on all NetBackup-supported platforms, the application server for the console does run on all supported platforms. The distributed application architecture enables direct administration of all NetBackup platforms, even though the consoles themselves run only on a subset of the NetBackup-supported platforms.

To log into the **NetBackup Administration Console**, specify a host name. The host name is the computer where the application server (`bpjava`) runs. (For example, a NetBackup primary server.) All requests or updates that are initiated in the console are sent to its application server that runs on this host.

Enhancing console performance

Performance of the NetBackup applications depends on the environment where the applications are running, including available resources and network throughput. The NetBackup default configuration, specifically the `INITIAL_MEMORY` and `MAX_MEMORY` configuration options, assumes sufficient memory resources on the computer where the console is running. For example, where the `jnbSA` command is run or the **NetBackup Administration Console** is started.

Following are guidelines for improving performance:

- Consider the network communication speed and the amount of data being transferred.
- Consider the amount of work being performed on the relevant computers. Run NetBackup on a computer that has a low level of activity. For example, there can be large differences in response time when other memory-intensive applications are running on the computer. (For example, web browsers.) Multiple instances of NetBackup on the same computer has the same effect.
- Run NetBackup on a 1-gigabyte computer that has at least 256 MB of RAM available to the application. In some instances, the application does not initiate due to insufficient memory. A number of messages identify these failures in the xterm window where the `jnbSA` command was run. Or, the messages appear in the application log file. Possible messages include the following:

```
Error occurred during initialization of VM
Could not reserve enough space for object heap
Out of Memory
```

See [“INITIAL_MEMORY, MAX_MEMORY”](#) on page 1069.

- Consider the amount of physical memory on the relevant computers. Possibly add memory on the host being administered (the console's application server host).
- Consider increasing the swap space to relevant computers:
 - The console host (the host where the console is started)
 - The host being administeredTo increase performance, increase the amount of swap space available to the system where you are running the applications. This is especially true if there is a great deal of other activity on the computer. More swap space can alleviate hangs or other problems that relate to insufficient memory for the applications.
- Consider additional or faster CPUs to relevant computers:
 - The console host (the host where the console is started)

- The host being administered
- Consider limiting the amount of NetBackup data that is retained for long periods of time to only that which is necessary. For example, do not retain successfully completed jobs for more than a few hours.
See [“About the jobs database”](#) on page 1032.

Determining better performance when the console is run locally or uses remote display back

Performance depends on the following:

- The speed of the network
- The console and the application server computer resources
- The workloads on the console
- The application server hosts
- The amount of NetBackup data (Data is the number of jobs in the Activity Monitor or number of NetBackup policies.)

The console may perform better if it is started on the console's application server host first, and then it is displayed back to the desktop host. However, little is known of a situation where that configuration produces better console performance. As previously mentioned, the configuration is not recommended due to problems unrelated to performance issues.

Consider the following scenarios to determine what would provide the best performance for your configuration.

NetBackup performance scenario 1

Assume no deficiency in either the console host's resources or the application server host's resources. Assume that the amount of NetBackup configuration data being transferred to the console host far exceeds the X-Windows pixel display data. That is, the actual console screen being sent from the remote host.

Unfortunately, the only way to determine the viability of this situation is to try it. Network capabilities and the proximity of the two hosts influences each NetBackup configuration.

NetBackup performance scenario 2

Assume that the available resources of the application server host far exceed that of the console host.

Assume that the console host has a very limited CPU and memory as compared to the NetBackup primary server being administered. (The console host is the

computer on which the console is started.) If the console is run on the primary server and displayed back to the desktop host, performance may be enhanced.

If the desktop host is a Windows computer, X-terminal emulation or remote display tools such as Exceed and VNC are required.

These scenarios address the performance aspect of using the NetBackup console. There may be other reasons that require you to display back remotely to your desktop, however, it is not recommended. Review the release notes for additional issues of relevance to the **NetBackup Administration Console** and the **Backup, Archive, and Restore** client console.

Table 33-2 shows the files that contain configuration entries.

Table 33-2 Files containing configuration entries

File	Description
/usr/opensv/java/auth.conf	Authorization options.
/usr/opensv/netbackup/bp.conf	Configuration options (server and client).
/usr/opensv/java/nbj.conf	Configuration options for the NetBackup Administration Console
/usr/opensv/volmgr/vm.conf	Configuration options for media and device management.
\$HOME/bp.conf	Configuration options for user (on client).

About adjusting time zones in the NetBackup Administration console

Sites in a geographically dispersed NetBackup configuration may need to adjust the time zone in the **NetBackup Administration Console** for administration of remote NetBackup hosts. (A remote NetBackup host may either be the host that is specified in the administration console logon dialog box or one referenced by the **File > Change Server** capability in the console.)

The default time zone for the console is that of the host on which the console is started, not the host that is specified (if different) in the console logon dialog box.

For backup, restore, or archive operations from within the **NetBackup Administration Console** or the **Backup, Archive, and Restore** application when run on a client, set the time zone relative to the NetBackup server from which the client restores files.

Set the time zone in separate instances of the **NetBackup Administration Console** when servers in different time zones are administered.

For example, open a **NetBackup Administration Console** to set the time zone for the local server in the Central time zone. To set the time zone for a server in the Pacific time zone as well, open another **NetBackup Administration Console**.

Change servers (**File > Change Server**), and then set the time zone for the Pacific time zone server. Doing so changes the time zone for the Central time zone server as well.

To perform a **Change Server** operation to a media server or client, the media server or client must have a security certificate installed.

Adjusting the time zone in the NetBackup Administration Console or the Backup, Archive, and Restore console

Use the following procedure to adjust the time zone or to use daylight savings time.

To adjust the time zone

- 1 In the **NetBackup Administration Console**, or in the **Backup, Archive, and Restore** console box, select **File > Adjust Application Time Zone**.
- 2 Select the **Standard** tab.
- 3 Clear the **Use custom time zone** check box.
- 4 Select the time zone.
- 5 For daylight savings time, select **Use daylight savings time**.
- 6 To have administrative capabilities and to apply the settings to the current session and all future sessions, select **Save as default time zone**.
- 7 Click **OK**.

Configuring a custom time zone in the NetBackup Administration Console or the Backup, Archive, and Restore console

Use the following procedure to configure a custom time zone in the administration or the client consoles.

To configure a custom time zone

- 1 In the **NetBackup Administration Console**, or in the **Backup, Archive, and Restore** console, select **File > Adjust Application Time Zone**.
- 2 Select the **Use custom time zone** check box.
- 3 Select the **Custom** tab.

- 4
- Select the time zone on which to base the **Backup, Archive, and Restore** interface time.
- 5
- For the **Offset from Greenwich Mean Time** setting, adjust the time to reflect how many hours and minutes the server's time zone is either behind or ahead of Greenwich Mean Time.
- 6
- Select the **Use daylight savings time** check box.
- 7
- To make a selection in the **Daylight savings time starts** section of the dialog, see the following table for descriptions of each option:

Begin daylight savings time on a specific date	Select Absolute date and indicate the month and day
Begin daylight savings time on the first occurrence of a day in a month	Select First day of week in month . Indicate the day of the week and the month.
Begin daylight savings time on the first occurrence of a day in a month and after a specific date	Select First day of week in month after date . Indicate the day of the week and the month and day.
Begin daylight savings time on the last occurrence of a day in a month	Select Last day of week in month . Indicate the day of the week and the month.
Begin daylight savings time on the last occurrence of a day in a month and before a specific date	Select Last day of week in month before date . Indicate the day of the week and the month and day.

- 8
- Indicate when daylight savings time should end by using one of the methods in the previous step.
- 9
- To have administrative capabilities and apply the settings to the current session and all future sessions, select **Save as default time zone**.
- 10
- Click **OK**.

Time zone table

This topic applies to NetBackup hosts that run the **NetBackup-Java Administration Console**. Sites in a geographically dispersed NetBackup configuration may need to adjust the time zone in the **NetBackup-Java Administration Console** for administration of remote NetBackup hosts.

The following table lists the available time zones in alphabetical order by ID name.

Table 33-3 Time zones

ID name	Display name
ACT	Central Standard Time (Northern Territory Australia)

Table 33-3 Time zones (*continued*)

ID name	Display name
AET	Eastern Standard Time (New South Wales)
Africa/Abidjan	Greenwich Mean Time
Africa/Accra	Greenwich Mean Time
Africa/Addis_Ababa	Eastern African Time
Africa/Algiers	Central European Time
Africa/Asmera	Eastern African Time
Africa/Bamako	Greenwich Mean Time
Africa/Bangui	Western African Time
Africa/Banjul	Greenwich Mean Time
Africa/Bissau	Greenwich Mean Time
Africa/Blantyre	Central African Time
Africa/Brazzaville	Western African Time
Africa/Bujumbura	Central African Time
Africa/Cairo	Eastern European Time
Africa/Casablanca	Western European Time
Africa/Ceuta	Central European Time
Africa/Conakry	Greenwich Mean Time
Africa/Dakar	Greenwich Mean Time
Africa/Dar_es_Salaam	Eastern African Time
Africa/Djibouti	Eastern African Time
Africa/Douala	Western African Time
Africa/El_Aaiun	Western European Time
Africa/Freetown	Greenwich Mean Time
Africa/Gaborone	Central African Time
Africa/Harare	Central African Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Africa/Johannesburg	South Africa Standard Time
Africa/Kampala	Eastern African Time
Africa/Khartoum	Eastern African Time
Africa/Kigali	Central African Time
Africa/Kinshasa	Western African Time
Africa/Lagos	Western African Time
Africa/Libreville	Western African Time
Africa/Lome	Greenwich Mean Time
Africa/Luanda	Western African Time
Africa/Lubumbashi	Central African Time
Africa/Lusaka	Central African Time
Africa/Malabo	Western African Time
Africa/Maputo	Central African Time
Africa/Maseru	South Africa Standard Time
Africa/Mbabane	South Africa Standard Time
Africa/Mogadishu	Eastern African Time
Africa/Monrovia	Greenwich Mean Time
Africa/Nairobi	Eastern African Time
Africa/Ndjamena	Western African Time
Africa/Niamey	Western African Time
Africa/Nouakchott	Greenwich Mean Time
Africa/Ouagadougou	Greenwich Mean Time
Africa/Porto-Novo	Western African Time
Africa/Sao_Tome	Greenwich Mean Time
Africa/Timbuktu	Greenwich Mean Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Africa/Tripoli	Eastern European Time
Africa/Tunis	Central European Time
Africa/Windhoek	Western African Time
AGT	Argentine Time
America/Adak	Hawaii-Aleutian Standard Time
America/Anchorage	Alaska Standard Time
America/Anguilla	Atlantic Standard Time
America/Antigua	Atlantic Standard Time
America/Araguaina	Brazil Time
America/Aruba	Atlantic Standard Time
America/Asuncion	Paraguay Time
America/Atka	Hawaii-Aleutian Standard Time
America/Barbados	Atlantic Standard Time
America/Belem	Brazil Time
America/Belize	Central Standard Time
America/Boa_Vista	Amazon Standard Time
America/Bogota	Colombia Time
America/Boise	Mountain Standard Time
America/Buenos_Aires	Argentine Time
America/Cambridge_Bay	Mountain Standard Time
America/Cancun	Central Standard Time
America/Caracas	Venezuela Time
America/Catamarca	Argentine Time
America/Cayenne	French Guiana Time
America/Cayman	Eastern Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
America/Chicago	Central Standard Time
America/Chihuahua	Mountain Standard Time
America/Cordoba	Argentine Time
America/Costa_Rica	Central Standard Time
America/Cuiaba	Amazon Standard Time
America/Curacao	Atlantic Standard Time
America/Danmarkshavn	Greenwich Mean Time
America/Dawson	Pacific Standard Time
America/Dawson_Creek	Mountain Standard Time
America/Denver	Mountain Standard Time
America/Detroit	Eastern Standard Time
America/Dominica	Atlantic Standard Time
America/Edmonton	Mountain Standard Time
America/Eirunepe	Acre Time
America/El_Salvador	Central Standard Time
America/Ensenada	Pacific Standard Time
America/Fort_Wayne	Eastern Standard Time
America/Fortaleza	Brazil Time
America/Glace_Bay	Atlantic Standard Time
America/Godthab	Western Greenland Time
America/Goose_Bay	Atlantic Standard Time
America/Grand_Turk	Eastern Standard Time
America/Grenada	Atlantic Standard Time
America/Guadeloupe	Atlantic Standard Time
America/Guatemala	Central Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
America/Guayaquil	Ecuador Time
America/Guyana	Guyana Time
America/Halifax	Atlantic Standard Time
America/Havana	Central Standard Time
America/Hermosillo	Mountain Standard Time
America/Indiana/Indianapolis	Eastern Standard Time
America/Indiana/Knox	Eastern Standard Time
America/Indiana/Marengo	Eastern Standard Time
America/Indiana/Vevay	Eastern Standard Time
America/Indianapolis	Eastern Standard Time
America/Inuvik	Mountain Standard Time
America/Iqaluit	Eastern Standard Time
America/Jamaica	Eastern Standard Time
America/Jujuy	Argentina Time
America/Juneau	Alaska Standard Time
America/Kentucky/Louisville	Eastern Standard Time
America/Kentucky/Monticello	Eastern Standard Time
America/Knox_IN	Eastern Standard Time
America/La_Paz	Bolivia Time
America/Lima	Peru Time
America/Los_Angeles	Pacific Standard Time
America/Louisville	Eastern Standard Time
America/Maceio	Brazil Time
America/Managua	Central Standard Time
America/Manaus	Amazon Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
America/Martinique	Atlantic Standard Time
America/Mazatlan	Mountain Standard Time
America/Mendoza	Argentine Time
America/Menominee	Central Standard Time
America/Merida	Central Standard Time
America/Mexico_City	Central Standard Time
America/Miquelon	Pierre and Miquelon Standard Time
America/Monterrey	Central Standard Time
America/Montevideo	Uruguay Time
America/Montreal	Eastern Standard Time
America/Montserrat	Atlantic Standard Time
America/Nassau	Eastern Standard Time
America/New_York	Eastern Standard Time
America/Nipigon	Eastern Standard Time
America/Nome	Alaska Standard Time
America/Noronha	Fernando de Noronha Time
America/North_Dakota/Center	Central Standard Time
America/Panama	Eastern Standard Time
America/Pangnirtung	Eastern Standard Time
America/Paramaribo	Suriname Time
America/Phoenix	Mountain Standard Time
America/Port_of_Spain	Atlantic Standard Time
America/Port-au-Prince	Eastern Standard Time
America/Porto_Acre	Acre Time
America/Porto_Velho	Amazon Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
America/Puerto_Rico	Atlantic Standard Time
America/Rainy_River	Central Standard Time
America/Rankin_Inlet	Eastern Standard Time
America/Recife	Brazil Time
America/Regina	Central Standard Time
America/Rio_Branco	Acre Time
America/Rosario	Argentine Time
America/Santiago	Chile Time
America/Santo_Domingo	Atlantic Standard Time
America/Sao_Paulo	Brazil Time
America/Scoresbysund	Eastern Greenland Time
America/Shiprock	Mountain Standard Time
America/St_Johns	Newfoundland Standard Time
America/St_Kitts	Atlantic Standard Time
America/St_Lucia	Atlantic Standard Time
America/St_Thomas	Atlantic Standard Time
America/St_Vincent	Atlantic Standard Time
America/Swift_Current	Central Standard Time
America/Tegucigalpa	Central Standard Time
America/Thule	Atlantic Standard Time
America/Thunder_Bay	Eastern Standard Time
America/Tijuana	Pacific Standard Time
America/Tortola	Atlantic Standard Time
America/Vancouver	Pacific Standard Time
America/Virgin	Atlantic Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
America/Whitehorse	Pacific Standard Time
America/Winnipeg	Central Standard Time
America/Yakutat	Alaska Standard Time
America/Yellowknife	Mountain Standard Time
Antarctica/Casey	Western Standard Time (Australia)
Antarctica/Davis	Davis Time
Antarctica/DumontDUrville	Dumont-d'Urville Time
Antarctica/Mawson	Mawson Time
Antarctica/McMurdo	New Zealand Standard Time
Antarctica/Palmer	Chile Time
Antarctica/South_Pole	New Zealand Standard Time
Antarctica/Syowa	Syowa Time
Antarctica/Vostok	Vostok Time
Arctic/Longyearbyen	Central European Time
ART	Eastern European Time
Asia/Aden	Arabia Standard Time
Asia/Almaty	Alma-Ata Time
Asia/Amman	Eastern European Time
Asia/Anadyr	Anadyr Time
Asia/Aqtau	Aqtau Time
Asia/Aqtobe	Aqtobe Time
Asia/Ashgabat	Turkmenistan Time
Asia/Ashkhabad	Turkmenistan Time
Asia/Baghdad	Arabia Standard Time
Asia/Bahrain	Arabia Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Asia/Baku	Azerbaijan Time
Asia/Bangkok	Indochina Time
Asia/Beirut	Eastern European Time
Asia/Bishkek	Kyrgyzstan Time
Asia/Brunei	Brunei Time
Asia/Calcutta	India Standard Time
Asia/Choibalsan	Choibalsan Time
Asia/Chongqing	China Standard Time
Asia/Chungking	China Standard Time
Asia/Colombo	Sri Lanka Time
Asia/Dacca	Bangladesh Time
Asia/Damascus	Eastern European Time
Asia/Dhaka	Bangladesh Time
Asia/Dili	East Timor Time
Asia/Dubai	Gulf Standard Time
Asia/Dushanbe	Tajikistan Time
Asia/Gaza	Eastern European Time
Asia/Harbin	China Standard Time
Asia/Hong_Kong	Hong Kong Time
Asia/Hovd	Hovd Time
Asia/Irkutsk	Irkutsk Time
Asia/Istanbul	Eastern European Time
Asia/Jakarta	West Indonesia Time
Asia/Jayapura	East Indonesia Time
Asia/Jerusalem	Israel Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Asia/Kabul	Afghanistan Time
Asia/Kamchatka	Petropavlovsk-Kamchatski Time
Asia/Karachi	Pakistan Time
Asia/Kashgar	China Standard Time
Asia/Katmandu	Nepal Time
Asia/Krasnoyarsk	Krasnoyarsk Time
Asia/Kuala_Lumpur	Malaysia Time
Asia/Kuching	Malaysia Time
Asia/Kuwait	Arabia Standard Time
Asia/Macao	China Standard Time
Asia/Macau	China Standard Time
Asia/Magadan	Magadan Time
Asia/Makassar	Central Indonesia Time
Asia/Manila	Philippines Time
Asia/Muscat	Gulf Standard Time
Asia/Nicosia	Eastern European Time
Asia/Novosibirsk	Novosibirsk Time
Asia/Omsk	Omsk Time
Asia/Oral	Oral Time
Asia/Phnom_Penh	Indochina Time
Asia/Pontianak	West Indonesia Time
Asia/Pyongyang	Korea Standard Time
Asia/Qatar	Arabia Standard Time
Asia/Qyzylorda	Qyzylorda Time
Asia/Rangoon	Myanmar Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Asia/Riyadh	Arabia Standard Time
Asia/Riyadh87	GMT+03:07
Asia/Riyadh88	GMT+03:07
Asia/Riyadh89	GMT+03:07
Asia/Saigon	Indochina Time
Asia/Sakhalin	Sakhalin Time
Asia/Samarkand	Turkmenistan Time
Asia/Seoul	Korea Standard Time
Asia/Shanghai	China Standard Time
Asia/Singapore	Singapore Time
Asia/Taipei	China Standard Time
Asia/Tashkent	Uzbekistan Time
Asia/Tbilisi	Georgia Time
Asia/Tehran	Iran Time
Asia/Tel_Aviv	Israel Standard Time
Asia/Thimbu	Bhutan Time
Asia/Thimphu	Bhutan Time
Asia/Tokyo	Japan Standard Time
Asia/Ujung_Pandang	Central Indonesia Time
Asia/Ulaanbaatar	Ulaanbaatar Time
Asia/Ulan_Bator	Ulaanbaatar Time
Asia/Urumqi	China Standard Time
Asia/Vientiane	Indochina Time
Asia/Vladivostok	Vladivostok Time
Asia/Yakutsk	Yakutsk Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Asia/Yekaterinburg	Yekaterinburg Time
Asia/Yerevan	Armenia Time
AST	Alaska Standard Time (United States)
Atlantic/Azores	Azores Time
Atlantic/Bermuda	Atlantic Standard Time
Atlantic/Canary	Western European Time
Atlantic/Cape_Verde	Cape Verde Time
Atlantic/Faeroe	Western European Time
Atlantic/Jan_Mayen	Eastern Greenland Time
Atlantic/Madeira	Western European Time
Atlantic/Reykjavik	Greenwich Mean Time
Atlantic/South_Georgia	South Georgia Standard Time
Atlantic/St_Helena	Greenwich Mean Time
Atlantic/Stanley	Falkland Island Time
Australia/ACT	Eastern Standard Time (New South Wales)
Australia/Adelaide	Central Standard Time (South Australia)
Australia/Brisbane	Eastern Standard Time (Queensland)
Australia/Broken_Hill	Central Standard Time (South Australia/New South Wales)
Australia/Canberra	Eastern Standard Time (New South Wales)
Australia/Darwin	Central Standard Time (Northern Territory)
Australia/Hobart	Eastern Standard Time (Tasmania)
Australia/LHI	Load Howe Standard Time
Australia/Lindeman	Eastern Standard Time (Queensland)
Australia/Lord_Howe	Load Howe Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Australia/Melbourne	Eastern Standard Time (Victoria)
Australia/North	Central Standard Time (Northern Territory)
Australia/NSW	Eastern Standard Time (New South Wales)
Australia/Perth	Western Standard Time (Australia)
Australia/Queensland	Eastern Standard Time (Queensland)
Australia/South	Central Standard Time (South Australia)
Australia/Sydney	Eastern Standard Time (New South Wales)
Australia/Tasmania	Eastern Standard Time (Tasmania)
Australia/Victoria	Eastern Standard Time (Victoria)
Australia/West	Western Standard Time (Australia)
Australia/Yancowinna	Central Standard Time (South Australia/New South Wales)
BET	Brazil Time
Brazil/Acre	Acre Time
Brazil/DeNoronha	Fernando de Noronha Time
Brazil/East	Brazil Time
Brazil/West	Amazon Standard Time
BST	Bangladesh Time
Canada/Atlantic	Atlantic Standard Time
Canada/Central	Central Standard Time
Canada/Eastern	Eastern Standard Time
Canada/East-Saskatchewan	Central Standard Time
Canada/Mountain	Mountain Standard Time
Canada/Newfoundland	Newfoundland Standard Time
Canada/Pacific	Pacific Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Canada/Saskatchewan	Central Standard Time
Canada/Yukon	Pacific Standard Time
CAT	Central African Time
CET	Central European Time
Chile/Continental	Chile Time
Chile/EasterIsland	Easter Island Time
CNT	Newfoundland Standard Time
CST	Central Standard Time (United States)
CST6CDT	Central Standard Time (United States)
CTT	China Standard Time
Cuba	Central Standard Time
EAT	Eastern African Time
ECT	Central European Time
EET	Eastern European Time
Egypt	Eastern European Time
Eire	Greenwich Mean Time
EST	Eastern Standard Time (United States)
EST5EDT	Eastern Standard Time (United States)
Etc/GMT	GMT+00:00
Etc/GMT	GMT+00:00
Etc/GMT+0	GMT+00:00
Etc/GMT+1	GMT-01:00
Etc/GMT+10	GMT-10:00
Etc/GMT+11	GMT-11:00
Etc/GMT+12	GMT-12:00

Table 33-3 Time zones (*continued*)

ID name	Display name
Etc/GMT+2	GMT-02:00
Etc/GMT+3	GMT-03:00
Etc/GMT+4	GMT-04:00
Etc/GMT+5	GMT-05:00
Etc/GMT+6	GMT-06:00
Etc/GMT+7	GMT-07:00
Etc/GMT+8	GMT-08:00
Etc/GMT+9	GMT-09:00
Etc/GMT-0	GMT-00:00
Etc/GMT-1	GMT+01:00
Etc/GMT-10	GMT+10:00
Etc/GMT-11	GMT+11:00
Etc/GMT-12	GMT+12:00
Etc/GMT-13	GMT+13:00
Etc/GMT-14	GMT+14:00
Etc/GMT-2	GMT+02:00
Etc/GMT-3	GMT+03:00
Etc/GMT-4	GMT+04:00
Etc/GMT-5	GMT+05:00
Etc/GMT-6	GMT+06:00
Etc/GMT-7	GMT+07:00
Etc/GMT-8	GMT+08:00
Etc/GMT-9	GMT+09:00
Etc/Greenwich	Greenwich Mean Time
Etc/UCT	Coordinated Universal Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Etc/Universal	Coordinated Universal Time
Etc/UTC	Coordinated Universal Time
Etc/Zulu	Coordinated Universal Time
Europe/Amsterdam	Central European Time
Europe/Andorra	Central European Time
Europe/Athens	Eastern European Time
Europe/Belfast	Greenwich Mean Time
Europe/Belgrade	Central European Time
Europe/Berlin	Central European Time
Europe/Bratislava	Central European Time
Europe/Brussels	Central European Time
Europe/Bucharest	Eastern European Time
Europe/Budapest	Central European Time
Europe/Chisinau	Eastern European Time
Europe/Copenhagen	Central European Time
Europe/Dublin	Greenwich Mean Time
Europe/Gibraltar	Central European Time
Europe/Helsinki	Eastern European Time
Europe/Istanbul	Eastern European Time
Europe/Kaliningrad	Eastern European Time
Europe/Kiev	Eastern European Time
Europe/Lisbon	Western European Time
Europe/Ljubljana	Central European Time
Europe/London	Greenwich Mean Time
Europe/Luxembourg	Central European Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Europe/Madrid	Central European Time
Europe/Malta	Central European Time
Europe/Minsk	Eastern European Time
Europe/Monaco	Central European Time
Europe/Moscow	Moscow Standard Time
Europe/Nicosia	Eastern European Time
Europe/Oslo	Central European Time
Europe/Paris	Central European Time
Europe/Prague	Central European Time
Europe/Riga	Eastern European Time
Europe/Rome	Central European Time
Europe/Samara	Samara Time
Europe/San_Marino	Central European Time
Europe/Sarajevo	Central European Time
Europe/Simferopol	Eastern European Time
Europe/Skopje	Central European Time
Europe/Sofia	Eastern European Time
Europe/Stockholm	Central European Time
Europe/Tallinn	Eastern European Time
Europe/Tirane	Central European Time
Europe/Tiraspol	Eastern European Time
Europe/Uzhgorod	Eastern European Time
Europe/Vaduz	Central European Time
Europe/Vatican	Central European Time
Europe/Vienna	Central European Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Europe/Vilnius	Eastern European Time
Europe/Warsaw	Central European Time
Europe/Zagreb	Central European Time
Europe/Zaporozhye	Eastern European Time
Europe/Zurich	Central European Time
GB	Greenwich Mean Time
GB-Eire	Greenwich Mean Time
GMT	Greenwich Mean Time
GMT0	GMT+00:00
Greenwich	Greenwich Mean Time
Hongkong	Hong Kong Time
HST	Hawaii Standard Time
Iceland	Greenwich Mean Time
IET	Eastern Standard Time
Indian/Antananarivo	Eastern African Time
Indian/Chagos	Indian Ocean Territory Time
Indian/Christmas	Christmas Island Time
Indian/Cocos	Cocos Islands Time
Indian/Comoro	Eastern African Time
Indian/Kerguelen	French Southern and Antarctic Lands Time
Indian/Mahe	Seychelles Time
Indian/Maldives	Maldives Time
Indian/Mauritius	Mauritius Time
Indian/Mayotte	Eastern African Time
Indian/Reunion	Reunion Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Iran	Iran Time
Israel	Israel Standard Time
IST	India Standard Time
Jamaica	Eastern Standard Time
Japan	Japan Standard Time
JST	Japan Standard Time
Kwajalein	Marshall Islands Time
Libya	Eastern European Time
MET	Middle Europe Time
Mexico/BajaNorte	Pacific Standard Time
Mexico/BajaSur	Mountain Standard Time
Mexico/General	Central Standard Time
Mideast/Riyadh87	GMT+03:07
Mideast/Riyadh88	GMT+03:07
Mideast/Riyadh89	GMT+03:07
MIT	West Samoa Time
MST	Mountain Standard Time (United States)
MST7MDT	Mountain Standard Time (United States)
Navajo	Mountain Standard Time (United States)
NET	Armenia Time
NST	New Zealand Standard Time
NZ	New Zealand Standard Time
NZ-CHAT	Chatham Standard Time
Pacific/Apia	West Samoa Time
Pacific/Auckland	New Zealand Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Pacific/Chatham	Chatham Standard Time
Pacific/Easter	Easter Island Time
Pacific/Efate	Vanuatu Time
Pacific/Enderbury	Phoenix Island Time
Pacific/Fakaofu	Tokelau Time
Pacific/Fiji	Fiji Time
Pacific/Funafuti	Tuvalu Time
Pacific/Galapagos	Galapagos Time
Pacific/Gambier	Gambier Time
Pacific/Guadalcanal	Solomon Island Time
Pacific/Guam	Chamorro Standard Time
Pacific/Honolulu	Hawaii Standard Time
Pacific/Johnston	Hawaii Standard Time
Pacific/Kiritimati	Line Island Time
Pacific/Kosrae	Kosrae Time
Pacific/Kwajalein	Marshall Islands Time
Pacific/Majuro	Marshall Islands Time
Pacific/Marquesas	Marquesas Time
Pacific/Midway	Samoa Standard Time
Pacific/Nauru	Nauru Time
Pacific/Niue	Niue Time
Pacific/Norfolk	Norfolk Time
Pacific/Noumea	New Caledonia Time
Pacific/Pago_Pago	Samoa Standard Time
Pacific/Palau	Palau Time

Table 33-3 Time zones (*continued*)

ID name	Display name
Pacific/Pitcairn	Pitcairn Standard Time
Pacific/Ponape	Ponape Time
Pacific/Port_Moresby	Papua New Guinea Time
Pacific/Rarotonga	Cook Island Time
Pacific/Saipan	Chamorro Standard Time
Pacific/Samoa	Samoa Standard Time
Pacific/Tahiti	Tahiti Time
Pacific/Tarawa	Gilbert Island Time
Pacific/Tongatapu	Tonga Time
Pacific/Truk	Truk Time
Pacific/Wake	Wake Time
Pacific/Wallis	Wallis and Futuna Time
Pacific/Yap	Yap Time
PLT	Pakistan Time
PNT	Mountain Standard Time
Poland	Central European Time
Portugal	Western European Time
PRC	China Standard Time
PRT	Atlantic Standard Time
PST	Pacific Standard Time (United States)
PST8PDT	Pacific Standard Time
ROK	Korea Standard Time
Singapore	Singapore Time
SST	Solomon Island Time
SystemV/AST4	Atlantic Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
SystemV/AST4ADT	Atlantic Standard Time
SystemV/CST6	Central Standard Time
SystemV/CST6CDT	Central Standard Time
SystemV/EST5	Eastern Standard Time
SystemV/EST5EDT	Eastern Standard Time
SystemV/HST1	Hawaii Standard Time
SystemV/MST7	Mountain Standard Time
SystemV/MST7MDT	Mountain Standard Time
SystemV/PST8	Pitcairn Standard Time
SystemV/PST8PDT	Pacific Standard Time
SystemV/YST9	Gambier Time
SystemV/YST9YDT	Alaska Standard Time
Turkey	Eastern European Time
UCT	Coordinated Universal Time
Universal	Coordinated Universal Time
US/Alaska	Alaska Standard Time
US/Aleutian	Hawaii-Aleutian Standard Time
US/Arizona	Mountain Standard Time
US/Central	Central Standard Time
US/Eastern	Eastern Standard Time
US/East-Indiana	Eastern Standard Time
US/Hawaii	Hawaii Standard Time
US/Indiana-Starke	Eastern Standard Time
US/Michigan	Eastern Standard Time
US/Mountain	Mountain Standard Time

Table 33-3 Time zones (*continued*)

ID name	Display name
US/Pacific	Pacific Standard Time
US/Pacific-New	Pacific Standard Time
US/Samoa	Samoa Standard Time
UTC	Coordinated Universal Time
VST	Indochina Time
WET	Western European Time
W-SU	Moscow Standard Time
Zulu	Coordinated Universal Time

Alternate server restores

This chapter includes the following topics:

- [About alternate server restores](#)
- [About supported configurations for alternate server restores](#)
- [About performing alternate server restores](#)

About alternate server restores

This topic explains how to restore files by using a server other than the one that was used to write the backup. This type of restore operation is called an alternate server restore or server independent restore. It allows easier access to data for restores in primary and media server clusters and provides better failover and disaster recovery capabilities.

The architecture of NetBackup allows storage devices to be located on multiple servers (either separate storage devices or a shared robot). The NetBackup image catalog on the primary server contains an entry that defines the server (primary or media server) to which each backup was written. Information specific to the backup media is contained within the primary server image catalog (in the attribute file for each backup). The information is also contained in the NetBackup database, located on the primary server.

To restore data through a device on another server is more involved than other restores. Use the methods that are described in this topic to restore the backups. Although the methods do not require you to expire and import backup images, in some instances it is useful.

The information in this topic is also pertinent in the case of restoring from a backup copy. If you created multiple copies of a backup, it is possible to restore from a specific backup copy other than the primary copy. To do so, use the `bprestore` command.


More information is available in the [NetBackup Commands Reference Guide](#).

About supported configurations for alternate server restores

[Figure 34-1](#) and [Figure 34-2](#) show configurations where NetBackup supports alternate server restores.

All methods for alternate server restores require that the server that is used for the restore be in the same cluster as the server that performed the original backup. The server must also share the same NetBackup database.


Figure 34-1 NetBackup servers that share robotic peripherals



Assume the following when NetBackup servers share robotic peripherals:

- A single, shared NetBackup database exists on the NetBackup primary server.
- The NetBackup primary server is available at the time of restore.
- Robotic control is on a NetBackup server that is available at the time of the restore.

Figure 34-2 NetBackup servers with separate non-shared peripherals



Assume the following when NetBackup have separate, non-shared robotic peripherals:

- The media is made physically accessible through an available NetBackup server. The NetBackup database is updated to reflect this move.
- A single, shared NetBackup database exists on the NetBackup primary server.
- The NetBackup primary server is available at the time of restore.
- Robotic control (if applicable) is on a NetBackup server that is available at the time of the restore.

About performing alternate server restores

The method that NetBackup administrators can use to perform alternate server restores depends on the configuration and the situation. The method can include one or more of the following:

- Modify the NetBackup catalogs.
See [“About modifying the NetBackup catalogs”](#) on page 1106.
- Override the original server for restores.
See [“Overriding the original server for restores”](#) on page 1107.
- Enable automatic failover to an alternate server

See [“About enabling automatic failover to an alternate server”](#) on page 1109.

About modifying the NetBackup catalogs

To perform alternate server restores by modifying the NetBackup catalogs, change the contents of the NetBackup catalogs. Use this method only when the server reassignment is permanent.

Some examples of when to use this method are as follows:

- Media is moved to an off-site location, where a media server exists.
- A robot was moved from one server to another.
- Two (or more) servers share a robot, each with connected drives and one of the servers is to be disconnected or replaced.
- Two (or more) servers each have their own robots. One of the server's robots has run out of media capacity for future backups, while several empty slots exist on another server's robot.

The actual steps that are used vary depending on whether the original server is still available.

Modifying NetBackup catalogs when the server that wrote the media is available

Use the following procedure to modify catalogs when the server that wrote the media is available.

To modify NetBackup catalogs when the server that wrote the media is available

- 1 If necessary, physically move the media.
- 2 Update the NetBackup database by using move volume options in the Media Manager administration utilities.
- 3 Update the NetBackup image catalog on the primary server.
- 4 Update the NetBackup media catalogs on both the original NetBackup server (`oldserver`) and the destination NetBackup server (`newserver`).

Use the following command, which can be run from any one of the NetBackup servers.

Enter the `admincmd` command on one line:

- As administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpmedia.exe -movedb -m media_id
-newserver hostname -oldserver hostname
```

- As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpmedia -movedb -m media_id -newserver hostname
-oldserver hostname
```

Modifying NetBackup catalogs when the server that wrote the media is unavailable

Use the following procedure to modify catalogs when the server that wrote the media is unavailable.

To modify NetBackup catalogs when the server that wrote the media is unavailable

- 1 If necessary, physically move the media.
- 2 Update the NetBackup database by using the move volume options.
- 3 Update only the NetBackup image catalog on the primary server.

Use the following commands from the NetBackup primary server.

Enter the `admincmd` command on one line:

- As administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpimage.exe -id media_id -newserver hostname
-oldserver hostname
```

- As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpimage -id media_id -newserver hostname
-oldserver hostname
```

Overriding the original server for restores

NetBackup allows the administrator to force restores to a specific server, regardless of where the files were backed up. For example, if files were backed up on server A, a restore request can be forced to use server B.

Examples of when to use this method are as follows:

- Two (or more) servers share a robot, each with connected drives. A restore is requested while one of the servers is either temporarily unavailable or is busy doing backups.
- A server was removed from the NetBackup configuration, and is no longer available.

Use the following procedure to override the original server for restores.

To override the original server for restores

- 1 Open the NetBackup web UI.
- 2 On the left, select **Hosts > Host properties**. Depending on the type of server to override, select the primary server or a media server. Then select **Edit primary server** or **Edit media server**.
 See [“General server properties”](#) on page 106.
- 3 Select **General server**.
- 4 Go to **Media host override** and select the **Add** button. Add entries for the original backup server and the restore server and select **Add**.
- 5 Select **Save**.

Overriding the original server for restores manually

Use the following procedure to manually override the original server for restores.

To manually override the original server for restores

- 1 If necessary, physically move the media and update the Enterprise Media Manager database Media Manager volume database to reflect the move.
- 2 Modify the NetBackup configuration on the primary server as follows:
 - By using the **NetBackup web UI**:
 On the left pane select **Hosts > Host properties**. Select the primary server and select **Edit primary server**. Open the **General server** host properties. Go to **Media host override** and select **Add**. Add entries for the original backup server and the restore server and select **Add**. Then select **Save**.
 - By modifying the `bp.conf` file on a UNIX NetBackup server:
 As `root` add the following entry to the `/usr/openv/netbackup/bp.conf` file:

```
FORCE_RESTORE_MEDIA_SERVER = fromhost tohost
```


The *fromhost* is the server that wrote the original backup and the *tohost* is the server to use for the restore.

To revert to the original configuration for future restores, delete the changes that were made in this step.

- 3 Stop and restart the NetBackup Request daemon on the primary server.

The override applies to all storage units on the original server. This means that restores for any storage unit on *fromhost* go to *tohost*.

About enabling automatic failover to an alternate server

NetBackup allows the administrator to configure automatic restore failover to an alternate server if the original server is temporarily inaccessible. Once it is configured, this method does not require administrator intervention.

See [“Restore failover properties”](#) on page 151.

Some examples of when to use this method are as follows:

- Two or more servers share a robot, each with connected drives.
When a restore is requested, one of the servers is temporarily inaccessible.
- Two or more servers have standalone drives of the same type.
When a restore is requested, one of the servers is temporarily inaccessible.

In these instances, inaccessible means that the connection between `bprd` on the primary server and `bptm` on the original server (through `bpcd`) fails.

Possible reasons for the failure are as follows:

- The original server is down.
- The original server is up but `bpcd` on that server does not respond. (For example, if the connection is refused or access is denied.)
- The original server is up and `bpcd` is fine, but `bptm` has problems. (For example, if `bptm` cannot find the required tape.)

Note: The failover uses only the failover hosts that are listed in the NetBackup configuration. By default, the list is empty and NetBackup does not perform the automatic failover.

Failing over to an alternate server

Use the following procedure to enable automatic failover to an alternate server.

To enable automatic failover to an alternate server

- 1 Modify the NetBackup configuration on the primary server as follows:
 - By using the **NetBackup web UI**:

On the left pane select **Hosts > Host properties**. Select the primary server and select **Edit primary server**. Open the **Restore failover** host properties. Select **Add**. Add entries for the media server and the failover restore servers and select **Add**. Then select **Save**.
 - By modifying the `bp.conf` file on a UNIX NetBackup server:

As `root`, add the following entry to the `/usr/openv/netbackup/bp.conf` file:

```
FAILOVER_RESTORE_MEDIA_SERVERS =  
failed_host host1 host2 ... hostN
```

Where:
The *failed_host* is the server that is not operational.
The *host1 ... hostN* are the servers that provide failover capabilities.

When automatic failover is necessary for a given server, NetBackup searches through the relevant `FAILOVER_RESTORE_MEDIA_SERVERS` list. NetBackup looks from left to right for the first server that is eligible to perform the restore.

There can be multiple `FAILOVER_RESTORE_MEDIA_SERVERS` entries and each entry can have multiple servers. However, a NetBackup server can be a *failed_host* in only one entry.
- 2 Stop and restart the NetBackup Request daemon on the primary server.

Expiring and importing media for alternate server restores

It may be necessary to expire media and then import it, even with the alternate server restore capabilities.

Regarding identifying media spanning groups, an alternate server restore operation can include media IDs that contain backup images that span media. It may be necessary to identify the media IDs that contain fragments of the spanned images. The group of related media is called a media spanning group.

On Windows: To identify the media in a specific media spanning group, run the following command as administrator from the command prompt on the NetBackup primary server:

```
cd install_path\NetBackup\bin  
bpimmedia.exe -spangroups -U -mediaid media_id
```

On UNIX: To identify the media in a specific media spanning group, run the following command as root on the NetBackup primary server:

```
cd /usr/opensv/netbackup/bin/admincmd  
bpimmedia -spangroups -U -mediaid media_id
```

To display all media in all spanning groups, omit `-mediaid media_id` from the command.

Managing client backups and restores

This chapter includes the following topics:

- [About server-directed restores](#)
- [About client-redirected restores](#)
- [About restoring the files that have Access Control Lists \(ACLs\)](#)
- [About setting the original atime for files during restores on UNIX](#)
- [Restoring the System State](#)
- [About the backup and restore of compressed files on VxFS file systems](#)
- [About backups and restores on ReFS](#)

About server-directed restores

A NetBackup user with the Administrator role or similar permissions can perform restores from the NetBackup primary server. This type of restore is available in the web UI for the following policy types:

BigData	Lotus-Notes	NBU-Catalog
Cloud-Object-Store	MongoDB Ops Manager	NDMP
Datastore	MS-Exchange-Server	Nutanix-AHV
FlashBackup	MS-SharePoint	Standard
FlashBackup-Windows	MS-Windows	Universal-Share

Hyper-V	MSDP-Object-Store	VMware (agent-based recovery)
Hypervisor – Nutanix	NAS-Data-Protection	
BigData	Hypervisor – Nutanix	NBU-Catalog
Cloud-Object-Store	Lotus-Notes	NDMP
Datastore		Nutanix-AHV
FlashBackup	MS-Windows	Standard
FlashBackup-Windows	MSDP-Object-Store	Universal-Share
Hyper-V	NAS-Data-Protection	VMware (agent-based recovery)

Other restore types are available for certain policy types.

Restore type	Supported policy typed
Archived backups	MS-Windows, Standard
Optimized backups	MS-Windows
Point-in-time rollback	MS-Windows, NAS-Data-Protection, Standard
Raw partition backups	FlashBackup, FlashBackup-Windows, Standard
True image backups	MS-Windows, NAS-Data-Protection, NBU-Catalog, Standard
Virtual disk restore	VMware
Virtual machine backups	Hyper-V, Hypervisor-Nutanix

Preventing server-directed restores for a client (Backup, Archive, and Restore interface)

By default, NetBackup clients are configured to allow NetBackup administrators on a primary server to direct restores to any client.

To prevent server-directed restores for a client do the following:

- On Windows clients:
Open the **Backup, Archive, and Restore** interface.

Select **File > NetBackup Client Properties > General**, then clear the **Allow server-directed restores** check box.

- On UNIX clients:
Add `DISALLOW_SERVER_FILE_WRITES` to the following file on the client:

```
/usr/opensv/netbackup/bp.conf
```

Note: On UNIX systems, the redirected restores can incorrectly set UIDs or GIDs that are too long. The UIDs and GIDs of files that are restored from one platform to another may be represented with more bits on the source system than on the destination system. If the UID or the GID name in question is not common to both systems, the original UID or GID may be invalid on the destination system. In this case, the UID or GID is replaced with the UID or GID of the user that performs the restore.

Generating progress logs on UNIX

On UNIX, no progress log is produced if the `bp.conf` file of the requesting server does not contain an entry for the restoring server. Without that entry, the restoring server has no access to write the log files to the requesting server. (A progress log is an entry in the **Task Progress** tab of the **Backup, Archive, and Restore** client interface.)

Consider the following solutions:

- To produce a progress log, add the requesting server to the server list.
Log on to the requesting server. In the NetBackup web UI, open the host properties for the primary server. Then select **Servers**. Add the restoring server to the server list.
- Log on to the restoring server. Go to the Activity monitor to determine the success of the restore operation.

To restore a UNIX backup that contains soft and hard links, run the **Backup, Archive, and Restore** client interface from a UNIX machine.

About client-redirected restores

The **Backup, Archive, and Restore** client interface contains options for allowing clients to restore the files that were backed up by other clients. The operation is called a redirected restore.

For the following Backup Services API (XBSA) agents, redirected restores to a different version of the agent is not supported:

- MariaDB
- MySQL
- PostgreSQL

If you use a non-root service user account, specific access must be allowed for that user when you add files to the `/usr/opensv/netbackup/db/altnames` directory. The service user account must have full access to these files through the ownership or group and the permissions. For example, if the service user is `svcname` and its group is `svrgrp`, the file can have permissions of `400`. If the file owner is for a different user and group, the file permissions must allow access to the service user. For example, `777`. Equivalent permission settings must be used in a Windows environment.

About restore restrictions

By default, NetBackup permits only the client that backs up files to restore those files. NetBackup ensures that the client name of the requesting client matches the peer name that was used to connect to the NetBackup server.

Unless clients share an IP address, the peer name is equivalent to the client's host name. (Clients can share an IP address due to the use of a gateway and token ring combination, or multiple connections.) When a client connects through a gateway, the gateway can use its own peer name to make the connection.

The NetBackup client name is normally the client's short host name, such as `client1` rather than a longer form such as `client1.null.com`.

The client name is found in the following location:

Open the **File > Backup, Archive, and Restore** interface. Click **File > Specify NetBackup Machines and Policy Type**. The client name that is selected as **Source client for restores** is the source of the backups to be restored.

Allowing all clients to perform redirected restores

The NetBackup administrator can allow clients to perform redirected restores. With this, all clients can restore the backups that belong to other clients.

To do so, first create an `altnames` directory on the NetBackup primary server where the backup policy for the clients resides. Place an empty `No.Restrictions` file inside of the directory.

- On Windows:
`install_path\NetBackup\db\altnames\No.Restrictions`
Do not add a suffix to the files in the `altnames` directory.

- On UNIX:

```
/usr/opensv/netbackup/db/altnames/No.Restrictions
```

The NetBackup client name setting on the requesting client must match the name of the client for which the backup was created. The peer name of the requesting client does not need to match the NetBackup client name setting.

Note: The `altnames` directory can present a breach of security, so use it only under limited circumstances. Users that are permitted to restore files from other clients may also have local permission to create the files that are found in the backup.

Caution: For security reasons, it is strongly advised not to use the `No.Restrictions` file approach. This approach enables any client to restore backups of other clients that can be a security threat.

Note: On using the `No.Restrictions` file approach, a notification is by default generated in the NetBackup web UI every 7 days. Use the `NOTIFY_SNOOZE_PERIOD_IN_DAYS` option to change the frequency of this notification from the default value to any value from 1 to 90.

See [“NOTIFY_SNOOZE_PERIOD_IN_DAYS option for NetBackup servers”](#) on page 283.

For information on alternative methods of alternate client restores, refer to the following topics:

See [“Allowing a single client to perform redirected restores”](#) on page 1116.

See [“Allowing redirected restores of a specific client’s files”](#) on page 1117.

Allowing a single client to perform redirected restores

The NetBackup administrator can permit a single client to restore the backups that belong to other clients.

To do so, create an `altnames` directory on the NetBackup primary server where the policy that backed up the other client(s) resides. Place an empty *peername* file inside of the `altnames` directory where *peername* is the client to possess restore privileges.

- On Windows:

```
install_path\NetBackup\db\altnames\peername
```

- On UNIX:

```
/usr/opensv/netbackup/db/altnames/peername
```


In this case, the requesting client (*peername*) can access the files that are backed up by another client. The NetBackup client name setting on *peername* must match the name of the other client.

Allowing redirected restores of a specific client's files

The NetBackup administrator can permit a single client to restore the backups that belong to another specific client.

To do so, create an `altnames` directory on the NetBackup primary server of the requesting client in the following location:

- On Windows:
`install_path\NetBackup\db\altnames\peername`
- On UNIX:
`/usr/opensv/netbackup/db/altnames/peername`

Then, create a *peername* file inside of the directory where *peername* is the client to possess restore privileges. Add to the *peername* file the names of the client(s) whose files the requesting client wants to restore.

The requesting client can restore the files that were backed up by another client if:

- The names of the other clients appear in the *peername* file, and
- The NetBackup client name of the requesting client is changed to match the name of the client whose files the requesting client wants to restore.

Examples of redirected restores

This topic provides some example configurations that allow clients to restore the files that were backed up by other clients. These methods may be required when a client connects through a gateway or has multiple Ethernet connections.

In all cases, the requesting client must have access to an image database directory on the primary server or the requesting client must be a member of an existing NetBackup policy.

- On Windows: `install_path\NetBackup\db\images\client_name`
- On UNIX: `/usr/opensv/netbackup/db/images/client_name`

Note: Not all file system types on all computers support the same features. Problems can be encountered when a file is restored from one file system type to another. For example, the S51K file system on an SCO computer does not support symbolic links nor does it support names greater than 14 characters long. You may want to restore a file to a computer that doesn't support all the features of the computer from which the restore was performed. In this case, all files may not be recovered.

In the following examples, assume the following conditions:

- *client1* is the client that requests the restore.
- *client2* is the client that created the backups that the requesting client wants to restore.
- On Windows: *install_path* is the path where you installed the NetBackup software. By default, this path is `C:\Program Files\Cohesity NetBackup`.

Note: On Windows: You must have the necessary permissions to perform the following steps.

On UNIX: You must be a root user for any of the steps that must be performed on the NetBackup server. You may also need to be a root user to make the changes on the client.

Example of a redirected client restore

Assume you must restore files to *client1* that were backed up from *client2*. The *client1* and *client2* names are those specified by the NetBackup client name setting on the clients.

On Windows:

- 1 Log on to the NetBackup server.
- 2 Add *client2* to the following file and perform one of the following:
 - Edit `install_path\NetBackup\db\altnames\client1` to include the name of *client2*.
 - Create the following empty file:
`install_path\NetBackup\db\altnames\No.Restrictions`

On UNIX:

- 1 Log on as root on the NetBackup server.
- 2 Perform one of the following actions:

- Edit `/usr/opensv/netbackup/db/altnames/client1` so it includes the name of *client2*. Or,
- Run the `touch` command on the following file:
`/usr/opensv/netbackup/db/altnames/No.Restrictions`

Note: The `No.Restrictions` file allows any client to restore files from *client2*.

- 3 Log on to *client1* and change the NetBackup client name to *client2*.
- 4 Restore the file.
- 5 Undo the changes that were made on the server and client.


Example of a redirected client restore using the altnames file

This example explains how `altnames` provides restore capabilities to clients that do not use their own host name when they connect to the NetBackup server.

By default, the NetBackup client name of the requesting client must match the peer name that is used in the connection to the NetBackup server. When the NetBackup client name is the host name for the client and matches the peer name (normal case), this requirement is met.

However, problems arise when clients connect to multiple ethernet or connect to the NetBackup server through a gateway.

Figure 35-1 Example restore from a token ring client



In this example, restore requests from *client1*, *client2*, and *client3* are routed through the TCP gateway. Because the gateway uses its own peer name rather than the client host names for connection to the NetBackup server, NetBackup refuses the requests. Clients cannot restore even their own files.

To correct the situation, do the following**1** Determine the peer name of the gateway:

- Try a restore from the client in question. In this example, the request fails with an error message similar to the following:

```
client is not validated to use the server
```

- Examine the NetBackup problems report and identify the peer name that is used on the request. Entries in the report may be similar to the following:

```
01/29/12 08:25:03 bpserver - request from invalid server or  
client client1.dvlp.null.com
```

In this example, the peer name is `client1.dvlp.null.com`.

2 Do one of the following:

On Windows: Determine the peer name, then create the following file on the NetBackup primary server:

```
install_path\NetBackup\db\altnames\peername
```

In this example, the file is:

```
install_path\NetBackup\db\altnames\client1.dvlp.null.com
```

On UNIX: Run the `touch` command on the following file:

```
/usr/opensv/netbackup/db/altnames/peername
```

In this example, the file is:

```
/usr/opensv/netbackup/db/altnames/client1.dvlp.null.com
```

3 Edit the *peername* file so that it includes the client names.

For example, if you leave file `client1.dvlp.null.com` empty, *client1*, *client2*, and *client3* can all access the backups that correspond to their NetBackup client name setting.

See [“Allowing a single client to perform redirected restores”](#) on page 1116.

If you add the names *client2* and *client3* to the file, you give these two clients access to NetBackup file restores, but exclude *client1*.

See [“Allowing redirected restores of a specific client’s files”](#) on page 1117.

Note that this example requires no changes on the clients.

4 Restore the files.

Example of how to troubleshoot a redirected client restore using the altnames file

If you cannot restore files with a redirected client restore by using the `altnames` file, troubleshoot the situation, as follows.

On Windows:

- 1 Create the debug log directory for the NetBackup Request Daemon:

```
install_path\NetBackup\logs\bprd
```

- 2 On the primary server, stop and restart the NetBackup Request Daemon. Restart the service to ensure that this service is running in verbose mode and logs information regarding client requests.
- 3 On *client1* (the requesting client), try the file restore.
- 4 On the primary server, identify the peer name connection that *client1* uses.
- 5 Examine the debug log for the NetBackup Request Daemon to identify the failing name combination:

Examine the failure as logged on the **All log entries** report. Or, examine the debug log for the NetBackup Request Daemon to identify the failing name combination:

```
install_path\NetBackup\logs\bprd\mmddyy.log
```

- 6 On the primary server, do one of the following:
 - Create an *install_path\NetBackup\db\altnames\No.Restrictions* file. The file allows any client to access *client2* backups if the client changes its NetBackup client name setting to *client2*.
 - Create an *install_path\NetBackup\db\altnames\peername* file. The file allows *client1* to access *client2* backups if *client1* changes its NetBackup client name setting to *client2*.
 - Add *client2* name to the following file:

```
install_path\NetBackup\db\altnames\peername.
```
 - *client1* is allowed to access backups on *client2* only.
- 7 On *client1*, change the NetBackup client name setting to match what is specified on *client2*.
- 8 Restore the files from *client1*.
- 9 Perform the following actions:
 - Delete *install_path\NetBackup\logs\bprd* and the contents.

- In the NetBackup web UI, open the host properties for the primary server. Click **Logging**. Clear the **Keep logs for days** setting.

10 If you do not want the change to be permanent, do the following:

- Delete `install_path\NetBackup\db\altnames\No.Restrictions` (if existent).
- Delete `install_path\NetBackup\db\altnames\peername` (if existent).
- On *client1*, change the NetBackup client name to its original value.

On UNIX:

1 On the NetBackup primary server, add the `VERBOSE` entry and a logging level to the `bp.conf` file. For example:

```
VERBOSE = 3
```

2 Create the debug log directory for `bprd` by running the following command:

```
mkdir /usr/opensv/netbackup/logs/bprd
```

3 On the NetBackup server, stop the NetBackup Request Daemon, `bprd`, and restart it in verbose mode by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdrege -terminate
/usr/opensv/netbackup/bin/bprd -verbose
```

Restart `bprd` to ensure that `bprd` logs information regarding client requests.

4 On *client1*, try the file restore.

5 On the NetBackup server, identify the peer name connection that *client1* used.

Examine the `bard debug` log to identify the failing name combination:

Examine the failure as logged on the **All log entries** report or examine the `bard debug` log to identify the failing name combination:

```
/usr/opensv/netbackup/logs/bprd/log.date
```

6 On the NetBackup server enter the following command:

```
mkdir -p /usr/opensv/netbackup/db/altnames touch
/usr/opensv/netbackup/db/altnames/No.Restrictions
```

This command allows any client access to *client2* backups by changing its NetBackup client name setting to specify the *client2*.

- 7 Run the touch command on the following file:

```
/usr/opensv/netbackup/db/altnames/peername
```

The command allows *client1* access to any *client2* backups by changing its NetBackup client name setting to specify *client2*.

- 8 Add *client2* to the `/usr/opensv/netbackup/db/altnames/peername` file. The addition to the *peername* file allows *client1* access to the backups that were created on *client2* only.
- 9 On *client1*, change the NetBackup client name setting in the user interface to match what is specified on *client2*.
- 10 Restore the files to *client1*.
- 11 Do the following:
 - Delete the `VERBOSE` entry from the `/usr/opensv/netbackup/bp.conf` file on the primary server.
 - Delete `/usr/opensv/netbackup/logs/bprd` and the contents.
- 12 Return the configuration to what it was before the restore.
 - Delete `/usr/opensv/netbackup/db/altnames/peer.or.hostname` (if it exists)
 - Delete `/usr/opensv/netbackup/db/altnames/No.Restrictions` (if it exists)
 - On *client1*, restore the NetBackup client name setting to its original value.

About restoring the files that have Access Control Lists (ACLs)

An Access Control List (ACL) is a table that conveys the access rights users need to a file or directory. Each file or directory can have a security attribute that extends or restricts users' access.

By default, the `nbtar` (`/usr/opensv/netbackup/bin/nbtar`) restores ACLs along with file and directory data.

However, in some situations the ACLs cannot be restored to the file data, as follows:

- Where the restore is cross-platform.
- When a restore utility (`tar`) other than `nbtar` is used to restore files.

In these instances, NetBackup stores the ACL information in a series of generated files in the `root` directory using the following naming form:

```
.SeCuRiT.y.nnnn
```

These files can be deleted or can be read and the ACLs regenerated by hand.

Note: If performing an alternate restore where the original directory was ACL-enabled, the alternate restore directory must also be ACL-enabled. If the alternate restore directory is not ACL-enabled, the restore is not successful.

Restore files without restoring ACLs

You can use the NetBackup web UI to restore data without restoring the ACLs. Both the destination client and the source of the backup must be Windows systems.

To restore files without restoring ACLs, the following conditions must be met:

- The policy that backed up the client is the policy type **MS-Windows**.
- The restore is performed from the server. The option is unavailable on standalone clients (clients that do not contain the NetBackup server software).
- You must have the RBAC role Administrator or a role with similar permissions. You must have the RBAC role **Administrator** or **Default Physical Clients** or have a role with similar permissions.
- The destination client and the source of the backup must both be systems running supported Windows OS levels. The option is disabled on UNIX clients.

Use the following procedure to restore files without restoring ACLs.

To restore files without restoring ACLs

- 1 Sign in to the primary server using the NetBackup web UI.
- 2 On the left, select **Recovery**.
- 3 On the **Regular recovery** card, select **Start recovery**.
- 4 Select the policy type **MS-Windows**.
- 5 Select the restore type **Normal backups**.
- 6 For the **Source client** select the client that was backed up.
- 7 (Optional) To redirect the restore to a different client, enter the **Destination client** name.
- 8 Select **Next**.
- 9 To restore files from a backup other than the most recent, follow these steps:
 You can restore from the most recent backup or select a specific backup:
 - If any backups for the client are within the specified date range, NetBackup populates the right pane with the information about the most recent backup. On the left, select the files.

- To restore files from a backup other than the most recent, follow these steps:
 - Next to **Date range** select **Edit**. Then select **Use backup history**.
 - Select the wanted backup and select **Apply**.
- 10 Select the files to restore. Then select **Next**.
- 11 In the **Recovery options**, select **Restore without access-control attributes**.
- 12 Make any other selections for the restore job. Then select **Next**.
- 13 Select **Start recovery**.

About setting the original atime for files during restores on UNIX

During a restore, NetBackup sets the `atime` for each file to the current time by default. You can elect to have NetBackup set the `atime` for each restored file to the value the file had when it was backed up. To do so, create the following file on the client:

```
/usr/openv/netbackup/RESTORE_ORIGINAL_ATIME
```

Restoring the System State

The System State includes the registry, the COM+ Class Registration database, and boot and system files. If the server is a domain controller, the data also includes the Active Directory services database and the SYSVOL directory.

Note: The best recovery procedure depends on many hardware and software variables that pertain to the server and its environment. For a complete Windows recovery procedure, refer to the Microsoft documentation.

Read the following notes carefully before you restore the System State:

- The System State should be restored in its entirety. Do not restore selected files.
- Do not redirect a System State restore. System State is computer-specific and to restore it to an alternate computer can result in an unusable system.
- Do not cancel a System State restore operation. To cancel the operation may leave the system unusable.

- To restore the System State to a domain controller, the Active Directory must not be running.

Restore the System State

Use the following procedure to restore the System State.

To restore the System State

- 1 To restore the Active Directory, restart the system, and press F8 during the boot process. F8 brings up a startup options menu. Press F8 upon restart if the system to which you are to restore is a Windows domain controller. Otherwise, begin with step 4.
- 2 From the startup options, select **Directory Services Restore Mode** and continue the boot process.
- 3 Ensure that the **NetBackup Client Service**, either `bpinetd` on Windows or `inetd` on UNIX, has started. Use the **Activity monitor** or the Windows Services application.
- 4 Sign in to the primary server using the NetBackup web UI.
- 5 On the left, select **Recovery**.
- 6 On the **Regular recovery** card, select **Start recovery**.
- 7 Select the policy type **MS-Windows**.
- 8 Select the restore type **Normal backups**.
- 9 For the **Source client** select the client that was backed up.
- 10 (Optional) To redirect the restore to a different client, enter the **Destination client** name.
- 11 Select **Next**.
- 12 To restore files from a backup other than the most recent, follow these steps:
You can restore from the most recent backup or select a specific backup:
 - If any backups for the client are within the specified date range, NetBackup populates the right pane with the information about the most recent backup. On the left, select the files.
 - To restore files from a backup other than the most recent, follow these steps:
 - Next to **Date range** select **Edit**. Then select **Use backup history**.
 - Select the wanted backup and select **Apply**.
- 13 Select the **System State**. Then select **Next**.

- 14** For the **Restore target**, select **Restore everything to its original location**.

Do not redirect the System State restore to a different host. System State is computer-specific. To restore it to a different computer can result in an unusable system.

- 15** In the Recovery options, select **Allow overwrite of existing files**.

- 16** Make any other selections for the restore job. Then select **Next**.

- 17** Select **Start recovery**.

- 18** The network may contain more than one domain controller. To replicate Active Directory to other domain controllers, perform an authoritative restore of the Active Directory after the NetBackup restore job completes.

To perform an authoritative restore of the Active Directory, run the Microsoft `ntdsutil` utility after you restored the System State data but before the server is restarted. An authoritative restore ensures that the data is replicated to all of the servers.

Additional information about an authoritative restore and the `ntdsutil` utility is available.

See the Microsoft documentation.

- 19** Restart the system before performing subsequent restore operations.

If you booted into **Directory Services Restore Mode** on a domain controller, restart into normal mode when the restore is complete.

About the backup and restore of compressed files on VxFS file systems

NetBackup can back up and restore VxFS-compressed files, maintaining the compression state when the target volume supports file system compression.

Upon backup of files on a VxFS file system, a message displays in the **Activity monitor** whenever NetBackup encounters a compressed file:

```
Compress flag found for 'file_name'.
```

Upon restore, NetBackup restores the files to a VxFS file system in their compressed form.

If the restore is to a non-VxFS file system, NetBackup restores the files in an uncompressed form. The following message displays in the **Details** tab for the job:

```
File 'file_name' will not be restored in compressed form.
```

The message appears only for the first file that cannot be restored in its compressed form.

Note: The compression messages display if the verbose level is 1 or greater.

About backups and restores on ReFS

Microsoft Resilient File System (ReFS) support in NetBackup is automatic and requires no additional configuration.

To restore optimized backups, use the **Restore from Optimized Backup** in the **Backup, Archive, and Restore** interface. You can select individual files for restore.

NetBackup does not support a redirected restore of a Microsoft Resilient File Systems (ReFS) file system.

[Table 35-1](#) lists the ReFS-to-NTFS backup and restore combinations and the success of each.

Table 35-1 ReFS backup and restore

Between file systems	Backups	Restores
ReFS to ReFS	Successful	Successful
ReFS to NTFS	Successful	Successful
NTFS to ReFS	Successful	Limited success For successful restores: - Restore NTFS backups to NTFS file system. - Remove all non-supported ReFS items. - Copy the files to an ReFS file system.

Known issue

A known issue exists that includes failures with respect to backups for files having ReFS based snapshot. At present Microsoft does not support backup of files having ReFS based snapshot as the API's are not compatible. Microsoft is working on documenting this behavior and providing support which are tracked with the following issue ID's:

- Documentation issue#: 42324557
- Backup Read issue#: 42295538

Powering down and rebooting NetBackup servers

This chapter includes the following topics:

- [Shut down and restart NetBackup servers](#)
- [Shut down and start up all NetBackup services and daemons](#)
- [Restart a NetBackup server](#)
- [Restart a NetBackup media server](#)

Shut down and restart NetBackup servers

To close and restart NetBackup servers, use the following recommended procedure.

To shut down a server

- 1 On the left, select **Activity monitor**. Select the **Jobs** tab and make sure that no jobs are running.
- 2 Select any daemon that is still running and select **Actions > Stop**.

- 3 From the command line, run:

On Windows:

```
install_path\NetBackup\bin\admincmd\bprdreq -terminate
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

`bprdreq` does not run on a media server.

- 4 Run the system shutdown command.

The installation process copies the appropriate startup and shutdown script from `/usr/opensv/netbackup/bin/goodies` to `/init.d` and creates links to it from the appropriate `/rc` directory.

Use system startup scripts to begin the Media Manager and NetBackup daemons when the system starts up. Use shutdown scripts to terminate the daemons at system shutdown.

The [NetBackup Installation Guide](#) contains more information about the startup scripts and shutdown scripts.

- 5 On Windows, run:

```
install_path\NetBackup\bin\bpdown
```

- 6 Shut down the server.

Shut down and start up all NetBackup services and daemons

To shut down and start all NetBackup services and daemons, enter the following commands from a command line:

On Windows:

- To shut down all NetBackup services:

```
install_path\NetBackup\bin\bpdown
```

- To startup all NetBackup services:

```
install_path\NetBackup\bin\bpup
```

On UNIX:

- To shut down all NetBackup daemons:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- To startup all NetBackup daemons:

```
/usr/opensv/netbackup/bin/bp.start_all
```

Restart a NetBackup server

Use the following procedure to restart a NetBackup server.

To restart a NetBackup primary server

- 1 Restart the system.
- 2 On Windows: If the required NetBackup services are not set up to start automatically, do the following:
 - From the Windows desktop, start the Windows Services applet.
 - Start the NetBackup Client service.
 - Start the NetBackup Device Manager service. The NetBackup Volume Manager service also starts automatically.
 - Start the NetBackup Request Daemon service to start the NetBackup Database Manager service.
- 3 On UNIX: Ensure that `bprd`, `bpdbm`, and `vmd` are up by running the following script:

```
/usr/opensv/netbackup/bin/bpps -a
```

- 4 On UNIX: Start all NetBackup daemons:

```
/usr/opensv/netbackup/bin/bp.start_all
```

Restart a NetBackup media server

Use the following procedure to restart a NetBackup media server.

To restart a NetBackup media server

- 1 Restart the system.
- 2 On Windows: The required NetBackup services start automatically if they are set up to do so.

If they are not set to start automatically, do the following:

- From the Windows desktop, start the Windows Services applet.
 - Start the NetBackup Client service.
 - Start the NetBackup Device Manager service (`ltid`). The NetBackup Volume Manager service (`vmg`) also starts.
- 3** On UNIX: Start `ltid` if it is not already running:
- On the left select **Activity monitor**, then select the **Processes** tab.
 - Select `ltid` and select **Start**.
- 4** On UNIX: From the command line, run:

```
/usr/opensv/volmgr/bin/ltid
```


About Granular Recovery Technology

This chapter includes the following topics:

- [About installing and configuring Network File System \(NFS\) for Active Directory Granular Recovery](#)
- [About configuring Services for Network File System \(NFS\)](#)
- [Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology \(GRT\)](#)
- [Configuring a different network port for NBFSD](#)

About installing and configuring Network File System (NFS) for Active Directory Granular Recovery

NetBackup uses Granular Recovery Technology (GRT) and Network File System (NFS) to recover the individual objects that reside within a database backup image, such as:

- A user account from an Active Directory database backup
- Email messages or folders from an Exchange database backup
- A document from a SharePoint database backup

The NetBackup client mounts and accesses a mapped drive over a secure connection to the NetBackup media server. The NetBackup media server handles the client requests through the NetBackup File System (NBFS) service, or NBFSD.

Multiple NetBackup agents that support GRT (for example, Exchange, SharePoint, and Active Directory) can use the same media server.

About configuring Services for Network File System (NFS)

To restore individual items from the Active Directory, you must configure Services for NFS on the NetBackup media server and all Active Directory domain controllers or ADAM/LDS hosts.

Table 37-1 Configuring NFS on Windows 2012, 2012 R2, or later


Step	Action	Description
Step 1	Configure NFS on the media server.	On the media server do the following: ■ Stop and disable the ONC/RPC Portmapper service, if it exists. ■ Enable NFS. See “Enabling Services for Network File System (NFS) on a media server” on page 1135. ■ Stop the Server for NFS service. See “Disabling the Server for NFS” on page 1142. ■ Stop the Client for NFS service. See “Disabling the Client for NFS on the media server” on page 1140. Note: If the Active Directory domain controller or ADAM/LDS host resides on the media server, do not disable the Client for NFS. ■ Configure the portmap service to start automatically at server restart. Issue the following from the command prompt: <code>sc config portmap start= auto</code> This command should return the status [SC] ChangeServiceConfig SUCCESS.
Step 2	Configure NFS on all Active Directory domain controllers or ADAM/LDS hosts.	On all Active Directory domain controllers or ADAM/LDS hosts, do the following: ■ Enable NFS on the clients. See “Enabling Services for Network File System (NFS) on a client” on page 1138. ■ Stop the Server for NFS service. See “Disabling the Server for NFS” on page 1142.

Enabling Services for Network File System (NFS) on a media server

To restore individual items from a backup that uses Granular Recovery Technology (GRT), you must enable Services for Network File System (NFS) on the media server. When this configuration is completed, you can disable any unnecessary NFS services.


To enable Services for Network File System (NFS) on a media server

- 1 Open the Server Manager.
- 2 From the **Manage** menu, click **Add Roles and Features**.




- 3 In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.


- 4
- On the **Select installation type** page, select **Role-based or feature-based installation**.




- 5
- Click **Next**.
- 6
- On the **Server Selection** page, click **Select a server from the server pool** and select the server. Click **Next**.



- 7 On the **Server Roles** page, expand **File and Storage Services** and **File and iSCSI Services**.
- 8 Click **File Server** and **Server for NFS**. When you are prompted, click **Add Features**. Click **Next**.



- 9 If the media server is also an Active Directory domain controllers or ADAM/LDS host, on the **Features** page, click **Client for NFS**. Click **Next**.




- 10** On the **Confirmation** page, click **Install**.
- 11** Disable any unnecessary services, as follows:
 - If you have a single host that functions as both the media server and the Active Directory domain controllers or ADAM/LDS host, you can disable the Server for NFS service.
 See [“Disabling the Server for NFS”](#) on page 1142.
 - For a host that is only the NetBackup media server, you can disable the Server for NFS and the Client for NFS services.
 See [“Disabling the Server for NFS”](#) on page 1142.
 See [“Disabling the Client for NFS on the media server”](#) on page 1140.

Enabling Services for Network File System (NFS) on a client


To restore individual items from a backup that uses Granular Recovery Technology (GRT), you must enable Services for Network File System (NFS). When this configuration is completed all the Active Directory domain controllers or ADAM/LDS hosts, you can disable any unnecessary NFS services.

To enable Services for Network File System (NFS) on a Windows client


- 1** Open the Server Manager.
- 2** From the **Manage** menu, click **Add Roles and Features**.




- 3 In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
- 4 On the **Select installation type** page, select **Role-based or feature-based installation**.



- 5 Click **Next**.
- 6 On the **Server Selection** page, click **Select a server from the server pool** and select the server. Click **Next**.



- 7 On the **Server Roles** page, click **Next**.
- 8 On the **Features** page, click **Client for NFS**. Click **Next**.



- 9 On the **Confirmation** page, click **Install**.


Disabling the Client for NFS on the media server

After you enable Services for Network File System (NFS) on a host that is only a NetBackup media server, you can disable the Client for NFS.

To disable the Client for NFS on the NetBackup media server


- 1 Open the Server Manager.
- 2 In the left pane, expand **Configuration**.

3 Click **Services**.



- 4 In the right pane, right-click on **Client for NFS** and click **Stop**.
- 5 In the right pane, right-click on **Client for NFS** and click **Properties**.

- 6 In the **Client for NFS Properties** dialog box, from the **Startup type** list, click **Disabled**.



- 7 Click **OK**.


Disabling the Server for NFS

After you enable Services for Network File System (NFS) on the media server and on the Active Directory domain controllers or ADAM/LDS hosts, you can disable Server for NFS.

To disable the Server for NFS


- 1 Open the Server Manager.
- 2 In the left pane, expand **Configuration**.

3 Click **Services**.



- 4 In the right pane, right-click on **Server for NFS** and click **Stop**.
- 5 In the right pane, right-click on **Server for NFS** and click **Properties**.

- 6 In the **Server for NFS Properties** dialog box, from the **Startup type** list, click **Disabled**.



- 7 Click **OK**.
- 8 Repeat this procedure for the media server and for all Active Directory domain controllers or ADAM/LDS hosts.

Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology (GRT)

To perform backups and restores that use Granular Recovery Technology (GRT), perform the following configuration if you use a UNIX media server and Windows clients:

- Confirm that your media server is installed on a platform that supports granular recovery.
For more information about supported platforms, see the *NetBackup Enterprise Server and Server - OS Software Compatibility List* at the following URL:

https://www.veritas.com/support/en_US/article.100040093

- No other configuration is required for the UNIX media server.
- Enable or install NFS on all Active Directory domain controllers or ADAM/LDS hosts.
 See “[Enabling Services for Network File System \(NFS\) on a media server](#)” on page 1135.
 See “[Enabling Services for Network File System \(NFS\) on a client](#)” on page 1138.
- You can configure a different network port for NBFSD.
 See “[Configuring a different network port for NBFSD](#)” on page 1145.

Configuring a different network port for NBFSD

NBFSD runs on port 7394. If another service uses the standard NBFSD port in your organization, you can configure the service on another port. The following procedures describe how to configure a NetBackup server to use a network port other than the default.

To configure a different network port for NBFSD (Windows server)

- 1 Log on as administrator on the computer where NetBackup server is installed.
- 2 Open Regedit.
- 3 Open the following key.:

`HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\Config`

- 4 Create a new DWORD value named **FSE_PORT**.
- 5 Right-click on the new value and click **Modify**.
- 6 In the **Value data** box, provide a port number between 1 and 65535.
- 7 Click **OK**.

To configure a different network port for NBFSD (UNIX server)

- 1 Log on as root on the computer where NetBackup server is installed.
- 2 Open the `bp.conf` file.
- 3 Add the following entry, where XXXX is an integer and is a port number between 1 and 65535.

`FSE_PORT = XXXX`